High-Frequency Monetary Policy Shocks: What Are Markets Learning?*

David Munro †

March 23, 2026

Abstract

Using tick-by-tick data for Federal Funds and Eurodollar futures, I document the high-frequency price discovery process following FOMC announcements. I find that price discovery often occurs almost instantaneously—frequently within seconds of the release—even for late-maturity futures, suggesting that mechanical arbitrage plays a significant role in determining futures prices around announcements. In contrast, announcements concerning unconventional policy, which typically require more intensive processing of qualitative language, exhibit a significantly more gradual price path—consistent with the idea that human interpretation operates on longer time scales. I reconstruct standard monetary policy shocks and show that shocks constructed using a narrow 15-second window are strongly correlated with those constructed over the conventional 20-minute window and carry similar economic content, further suggesting that mechanical price adjustment—rather than the gradual learning about policy information often assumed in the literature—accounts for much of what these shocks identify. If mechanical arbitrage shapes shock magnitudes, frictions that impede arbitrage should matter. Consistent with this, I find that measures of financial market stress (prior to announcements) explain a substantial share of the cross-meeting variation in shock magnitudes, with funding market conditions most important for target-rate shocks and interest rate uncertainty most important for path-factor shocks.

Keywords: monetary policy shocks, FOMC announcements, high frequency
JEL Codes: E52, E43, G12, C58

*Preliminary draft, comments and suggestions welcome. I thank participants at the 2026 AALAC Workshop on Monetary Economics in the Digital Era at Grinnell College for helpful comments and suggestions, and Asher Jenkins and Xianzhi Wang for excellent research assistance. Funding from Middlebury College is gratefully acknowledged.

†Department of Economics, Middlebury College *dmunro@middlebury.edu*

1 Introduction

High-frequency identification has become a central tool for measuring the causal effects of monetary policy. By focusing on asset price movements in a narrow window around Federal Open Market Committee (FOMC) announcements, this approach minimizes contamination from other macroeconomic news and delivers shocks that are widely interpreted as exogenous innovations to monetary policy. These high-frequency monetary policy shocks now play a foundational role in empirical macroeconomics and finance.

A growing literature argues that these shocks capture more than just unexpected changes in the current policy rate. FOMC announcements may reveal information about the future path of policy (e.g. Gürkaynak et al. (2005) and Bauer and Swanson (2023)), the Federal Reserve's private information about the economy (Nakamura and Steinsson, 2018), its assessment of risks, or its policy reaction function. Under this view, high-frequency asset price responses partly reflect market participants interpreting the language and signals contained in FOMC communications. As a result, the shocks identified in standard high-frequency windows are often decomposed into "target" and "path" factors, or interpreted as containing a central bank information component.

A key but largely unexamined assumption underlying these interpretations concerns *who* is processing this information and *how quickly* they can do so. The standard identification window measures price changes from 10 minutes before to 20 minutes after the announcement (the "10/20 window"). This window is short enough to limit contamination from other news, yet long enough to plausibly allow human traders to read, interpret, and trade on the information contained in the announcement. Implicit in many interpretations of high-frequency shocks is the idea that at least part of the observed price adjustment reflects this human processing of complex policy communication.

In this paper, I argue that the timing of market responses calls this interpretation into question. Using trade-level data for Federal Funds and Eurodollar futures, I document that asset prices adjust to many FOMC announcements *extremely* quickly, with a substantial fraction of the total 10/20 window response occurring within the first few seconds of the

release. This pattern is especially stark for the current-month federal funds futures contract, whose payoff depends on the contemporaneous policy rate. However, I also show that longer-maturity Eurodollar contracts—often interpreted as reflecting forward guidance or policy path information—frequently adjust at comparable speeds. These dynamics are difficult to reconcile with a view in which traders gradually incorporate nuanced information from the text of the FOMC statement. Instead, the evidence is more consistent with algorithmic trading and rapid cross-market arbitrage triggered by machine-readable information releases.

This interpretation is consistent with the structure of modern financial markets. By the mid-2010s, an estimated 80 percent of trading volume in U.S. futures markets originated from high-frequency trading firms (Miller and Shorter, 2016). In an environment dominated by algorithmic order flow, it is perhaps unsurprising that price discovery occurs at machine speed. Crucially, if mechanical arbitrage is a primary driver of price adjustment, then conditions in financial markets that impede arbitrage should have measurable effects on the magnitude of the resulting shocks. The second part of the paper explores this implication directly.

To provide a sharper benchmark for interpretation, I examine episodes of unconventional policy, including large-scale asset purchase announcements and explicit forward guidance at the effective lower bound. Unlike conventional rate decisions, these announcements are text-heavy and require economic interpretation; they do not map mechanically into asset prices. Consistent with this distinction, price discovery during these episodes unfolds more gradually, often over 20 to 30 minutes. This contrast suggests that when interpretation and economic reasoning are required, adjustment occurs on time scales plausibly consistent with human processing.

While these event studies are informative, they are necessarily selective. To assess the broader implications for the literature, I replicate two widely used approaches to constructing high-frequency monetary policy shocks. First, following Nakamura and Steinsson (2018), I extract the first principal component from changes in current- and near-term interest rate futures. I show that shocks constructed using an ultra-narrow 15-second window are highly

correlated with shocks constructed using the conventional 20-minute window. Moreover, the two series exhibit nearly identical relationships with Treasury yields across maturities. These findings cast doubt on the interpretation that the conventional shock primarily reflects a slower process of human learning about forward guidance or central bank information. If such learning were central, shocks measured over 15 seconds—too short a window for textual interpretation—should differ meaningfully from their 20-minute counterparts.

Second, I replicate the target and path factor decomposition of Gürkaynak et al. (2005). Here too, the path factor estimated using a 15-second window is strongly correlated with its 20-minute counterpart and displays similar effects on the yield curve, implying that a substantial portion of what is labeled "path" information is incorporated within seconds of the announcement.

Together, these findings suggest that mechanical arbitrage plays a significant role in shaping futures prices around FOMC announcements. This motivates a natural question: if arbitrage is operative, do frictions that limit arbitrage affect the magnitude of the identified shocks? A substantial body of work in finance documents that arbitrage relationships break down during periods of financial market stress, when intermediary balance sheet constraints tighten (e.g., Du et al. (2018)). I bring this perspective to the construction of high-frequency monetary policy shocks. Using three complementary indicators of financial market conditions—the ICE BofA MOVE Index, the Hu et al. (2013) Treasury illiquidity measure, and the TED spread—I show that these measures explain a substantial share of the cross-meeting variation in the magnitude of both target-rate and path-factor shocks. Strikingly, the relevant stress measure differs across shock types: funding market stress (TED spread) is the dominant predictor of target-rate shock magnitude, while interest rate uncertainty (MOVE) better explains variation in path-factor shocks. This distinction is consistent with the different arbitrage mechanisms underlying each factor.

These findings may offer a new perspective on why high-frequency monetary policy shocks sometimes generate puzzling macroeconomic effects (e.g. Acosta (2022)). If shocks partly reflect mechanical arbitrage rather than deliberative learning, their informational content may

differ from what is typically assumed—and their magnitude may be sensitive to prevailing market conditions in ways that are difficult to interpret through the lens of policy surprises alone.

Collectively, the evidence suggests that understanding the microstructure and timing of price discovery is essential for interpreting high-frequency monetary policy shocks. The speed at which markets incorporate announcement information—and the sensitivity of price adjustments to financial market conditions—calls for a closer examination of the microstructure foundations of these shocks before attributing them to news about current and future monetary policy.

2 High-Frequency Methodology and Data

2.1 High-Frequency Identification of Monetary Policy Shocks

The high-frequency identification approach measures monetary policy shocks as changes in asset prices in a narrow window around policy announcements. Futures contracts are particularly well suited for this purpose because their prices embed market expectations of future policy rates. In an efficient market, discrete changes in these prices in direct response to FOMC announcements reflect new information about monetary policy.

To measure surprises to the current policy rate, the literature typically uses Federal Funds Futures (FFF). These are monthly contracts whose settlement value is determined by the arithmetic average of the effective federal funds rate over the delivery month. As a result, FFF prices provide a direct market-based measure of expectations about the near-term level of the policy rate, and changes in these prices in a narrow window around FOMC announcements are interpreted as unexpected revisions to the current policy stance.

To capture changes in expectations about the future path of policy, researchers often turn to longer-maturity interest rate futures, most commonly two- to four-quarter-ahead Eurodollar futures. These contracts are based on the 3-month U.S. dollar London Interbank Offered Rate (LIBOR), which historically served as a benchmark for short-term dollar funding costs. Although LIBOR has since been phased out and replaced by alternative reference rates

(SOFR), Eurodollar futures were the dominant instrument over most of the sample period for inferring expectations of short-term interest rates several quarters ahead. Movements in these longer-maturity contracts around FOMC announcements are therefore interpreted as revisions to the expected trajectory of future short-term rates, beyond the current month.

Following much of the literature, price changes are typically measured from 10 minutes before to 20 minutes after the announcement (the "10/20" window). I adopt this convention as a benchmark before exploring the sensitivity of measured shocks to alternative, narrower windows.

In practice, many interest rate futures contracts—particularly earlier in the sample—trade infrequently. As a result, high-frequency shocks are typically constructed using the last observed transaction price prior to the start and end of the identification window. While this "last price" approach guarantees a well-defined price change for each meeting, it abstracts from the intraday path of prices and obscures the timing of adjustment within the window. In particular, a 10/20 window shock does not reveal whether prices moved immediately at the announcement and then stabilized or adjusted gradually over many minutes.

Distinguishing between these possibilities is central to interpreting what high-frequency shocks represent. If prices adjust within seconds of the announcement, this is more consistent with rapid algorithmic trading and mechanical cross-contract arbitrage. If instead prices evolve gradually over several minutes, this would be more consistent with human traders processing textual information and revising beliefs about the economic outlook or the future path of policy. A central objective of this paper is therefore to provide more insight on within-window price changes and examine the second-by-second evolution of prices.

2.2 Data

I use tick-by-tick transaction data for Federal Funds and Eurodollar futures traded on the Chicago Mercantile Exchange (CME). The sample spans scheduled FOMC announcements from 1997 to 2023.[1] Each observation records the transaction time and price, allowing me

[1]Eurodollar futures were discontinued in 2023 with the end of LIBOR. But researchers have extended same approaches to constructing monetary policy shocks using SOFR futures Acosta et al. (2024).

to construct high-frequency price series around announcement times.

Accurate announcement timestamps are crucial for analyzing price dynamics at very high frequencies. For earlier years in the sample, I use announcement times provided in existing high-frequency monetary policy datasets used in the literature (Gürkaynak et al. (2005) and Nakamura and Steinsson (2018)). For later years, I update these timestamps using release times from Bloomberg Newswire headlines associated with FOMC policy announcements. Bloomberg timestamps record the moment at which policy decisions and statement headlines are disseminated through machine-readable news feeds.

While no timestamp perfectly captures the exact instant at which information first becomes available to all market participants, Bloomberg newswire times are particularly relevant for this study. These feeds are widely used by market participants and are directly integrated into many algorithmic trading systems. As a result, Bloomberg timestamps closely reflect the timing of information arrival to the automated trading infrastructure that plays a central role in modern financial markets.

These timestamps should be interpreted as measuring when policy information becomes available to the electronic news infrastructure used by many market participants, rather than the precise instant at which information first exists in any form. In modern markets, algorithmic trading systems frequently use structured newswire feeds directly, making these timestamps economically meaningful markers for the arrival of tradable information. However, small discrepancies may arise if some traders obtain information through alternative channels slightly before or after the newswire release. As a result, the timestamps used here provide a close, but not necessarily perfect, proxy for the timing of information arrival, which is important to keep in mind when interpreting reaction speeds measured in seconds.

For each meeting, I align transactions in event time relative to the announcement timestamp and construct second-by-second price series using a previous-tick interpolation method. This approach produces step-function price paths that accurately reflect the timing of price discovery in markets with irregular trading. These data allow me to analyze how quickly prices adjust to monetary policy announcements and how this adjustment varies across con-

tracts and over time.

3 Results

3.1 Visualizing changes around FOMC announcements

Before turning to the intraday price dynamics around FOMC announcements, it is useful to first document the overall magnitude of announcement effects. Following the standard 10/20 window approach, I compute the change in the implied Federal Funds rate from Fed Funds Futures (FFF) for each meeting and plot the empirical cumulative distribution function (ECDF) of these shocks in Figure 1.

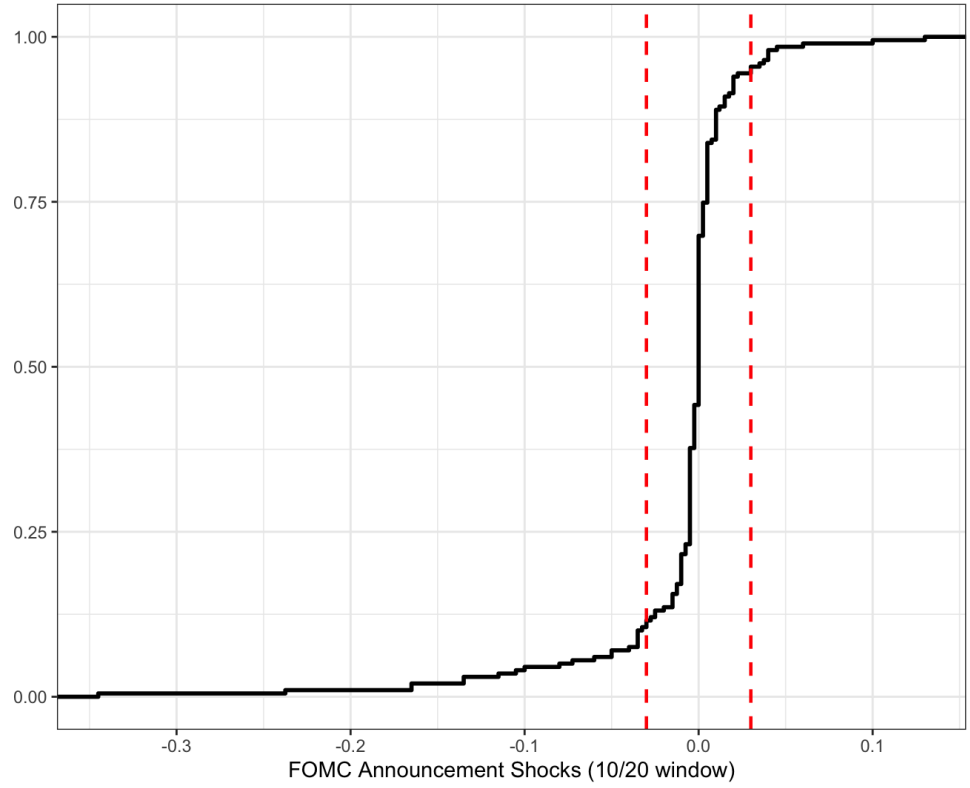


Figure 1: FOMC Announcement Changes in Federal Fund Futures Implied Interest Rate

Figure 1 shows that most announcement-induced changes are quite small. This is consistent with the view that FOMC communications and forward guidance often allow markets

to anticipate policy decisions well in advance of the meeting, limiting the degree of surprise at the time of the formal announcement.

To study the high-frequency dynamics of market reactions, I focus on meetings associated with economically meaningful surprises — small/no market reactions would naturally not show meaningful dynamics since markets are not responding to the announcement. For the examples of intra-day dynamics I present below, I select shocks with absolute changes exceeding 3 basis points. These thresholds are plotted as dashed red lines in 1, and highlight that shocks of these magnitudes are relatively infrequent, representing less than 20% of all announcements.

3.1.1 Intraday Price Dynamics

Below are four examples of the dynamics of the implied FFR around FOMC announcements, the timing of these releases are depicted with the dashed vertical red line. To give a broad sense of the temporal evolution of implied interest rates around announcements, I plot 2 hours before and after the announcement release.

These examples illuminate a few interesting characteristics. The first characteristic that it is that liquidity in the Fed Fund Futures market has changed dramatically over time.The top two panels of Figure 2 are large shocks from 1997 and 1999, respectively. In both cases, there is a small number of trades in the 4 hour (centered) window around the announcement, and in both cases there are hour-long windows with no trades. Comparing these first two announcements to the large example shocks in 2007 and 2019, we see a dramatic increase in the frequency of trading activity.

The second observation from these example shocks is that the changes in the implied FFR manifest themselves very quickly. To emphasize the speed of these market adjustments, the bottom two figures are re-plotted for a 10 minute window spanning the announcement release.

Zooming into the ten minutes surrounding the announcement (five minutes before to five minutes after) highlights that markets respond to FOMC releases *extremely* quickly.[2] This

[2]In Appendix A.1 I plot volume dynamics for these two examples, along with the examples presented for the Eurodollar contracts below. These also show a notable spike in the volume of trading immediately

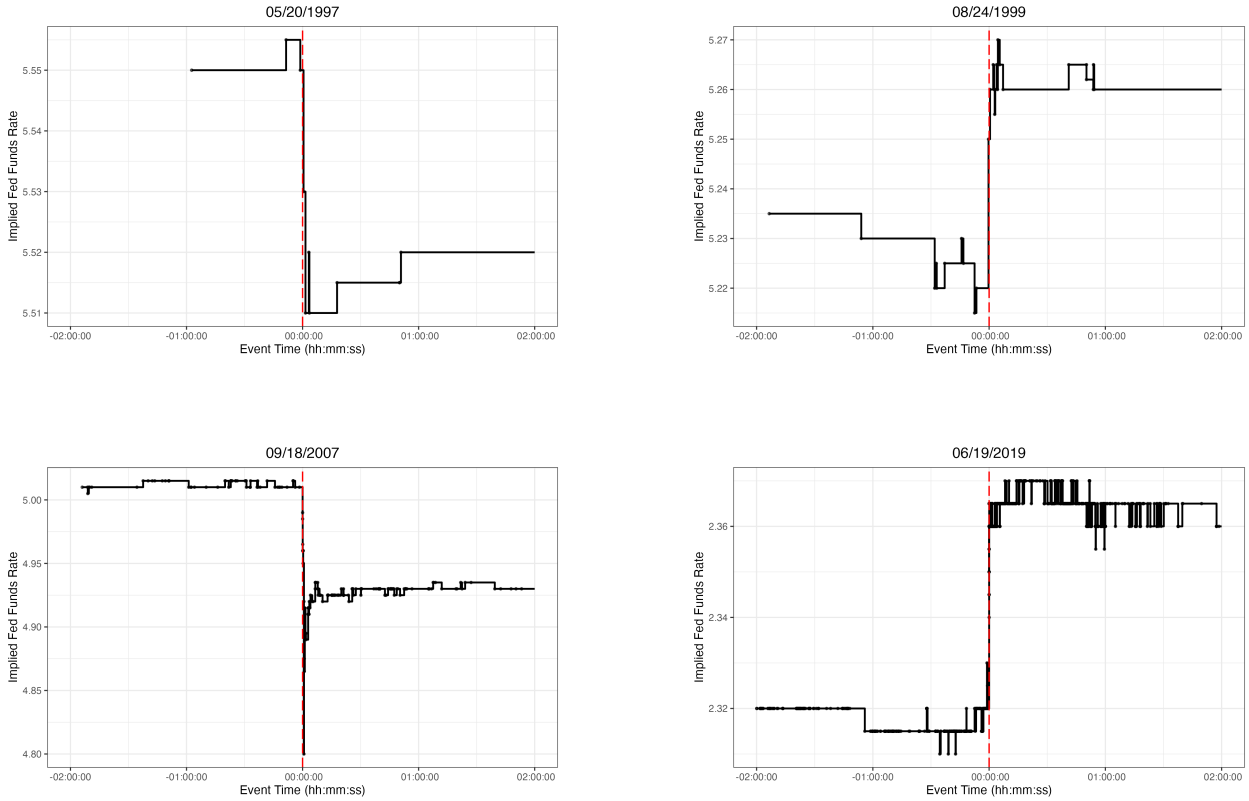


Figure 2: Four Examples of Large Shocks in Fed Fund Futures (+/- 2hr windows)

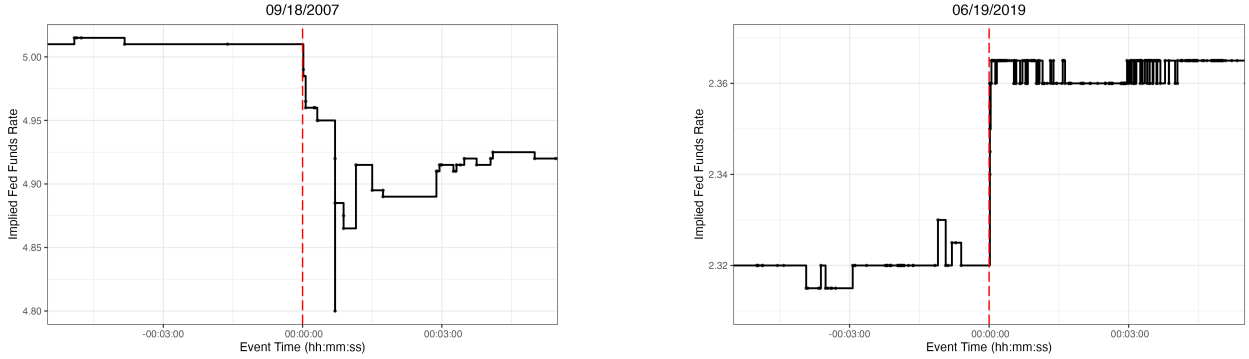


Figure 3: Example large shocks in 10 minute window

rapid adjustment is perhaps unsurprising for the current-month federal funds futures contract. By construction, the payoff of this contract depends directly on the realized effective federal funds rate over the remainder of the delivery month. As noted by Paul (2020), the only component of the FOMC announcement that is mechanically relevant for pricing this contract is the contemporaneous policy rate decision itself. Information about the Federal Reserve's outlook for the economy, its assessment of risks, or its intentions for policy beyond the current meeting should, in principle, matter primarily for contracts expiring at later horizons.

This institutional feature provides a natural interpretation of the speed of adjustment observed in the data. Once the policy rate decision is revealed, the uncertainty relevant for pricing the current-month contract is largely resolved, and traders can quickly arbitrage prices toward their new equilibrium. Whether these trades are executed by algorithms or by human traders reacting to a clearly stated numerical policy decision is secondary. The key point is that the information required to reprice the contract is simple, discrete, and immediately observable. The near-instantaneous response of the current-month federal funds futures contract is therefore consistent with trading on mechanically interpretable policy rate news, rather than with slower, interpretive learning about broader economic information embedded in the announcement text.

Importantly, this logic does not extend as cleanly to longer-maturity contracts. Futures on subsequent months and on two- to four-quarter-ahead Eurodollar contracts embed expectations about the future path of short-term interest rates and therefore should be sensitive to forward guidance and to information about the Federal Reserve's assessment of the economy. Mapping textual guidance or signals about economic conditions into revisions of expected future short rates requires more interpretation and economic reasoning than simply observing the new policy rate. If financial markets are actively learning about such information from FOMC announcements, one would expect price adjustment in these longer-maturity contracts to unfold more gradually than in the current-month contract. The next section

after FOMC announcements, particularly for Eurodollar futures.

examines this hypothesis by comparing the temporal dynamics of price responses across maturities.

3.1.2 Eurodollar Futures

Price dynamics in the FFF data presented above primarily reflect surprises to the current policy rate, as future policy decisions lie outside the duration of the 1-month contract. Because the federal funds rate is readily arbitrageable and heavily monitored by algorithmic traders, surprises to the current target should be incorporated into prices almost immediately.

To study how policy announcements affect expectations of the future path of policy, researchers turn to longer-dated instruments. The prices (implied interest rates) of later-expiration futures embed both the current policy rate and expectations of future policy actions, consistent with the expectations hypothesis of the term structure. A surprise change in the current policy rate mechanically affects these contracts through arbitrage relationships, leading to rapid adjustment via calendar spread trading and cross-market arbitrage (e.g., relative to OIS swaps).

By contrast, revisions to expectations about the future path of policy — particularly those arising from qualitative language in the announcement — require interpretation. Such information may diffuse more gradually as market participants process forward guidance, reassess the reaction function, and update beliefs.

The core argument of this paper is that these two channels — mechanical repricing of the current rate versus belief updating about future policy — imply different intraday adjustment dynamics. Rapid, near-instantaneous movements in longer-dated futures are more likely to reflect arbitrage-based transmission of the current rate surprise, whereas more gradual price evolution may reflect learning from the informational content of the announcement language. Accordingly, the second-by-second dynamics of Eurodollar futures provide insight into the underlying source of price changes.

Following the general approach in Brennan et al. (2025), I examine the FFF contract for the next scheduled FOMC meeting, as well as Eurodollar futures expiring 2, 3, and 4

quarters ahead. Using the notation in Brennan et al. (2025), shocks constructed from the current-month and next-meeting FFF contracts are denoted MP1 and MP2, respectively, while shocks constructed from the Eurodollar contracts are denoted ED2, ED3, and ED4 according to their horizon. Figure 4 plots the raw second-by-second price series in 10 minute (centered) windows for MP1 through ED4 for four FOMC announcements, illustrating how the speed and persistence of price discovery vary across maturities.

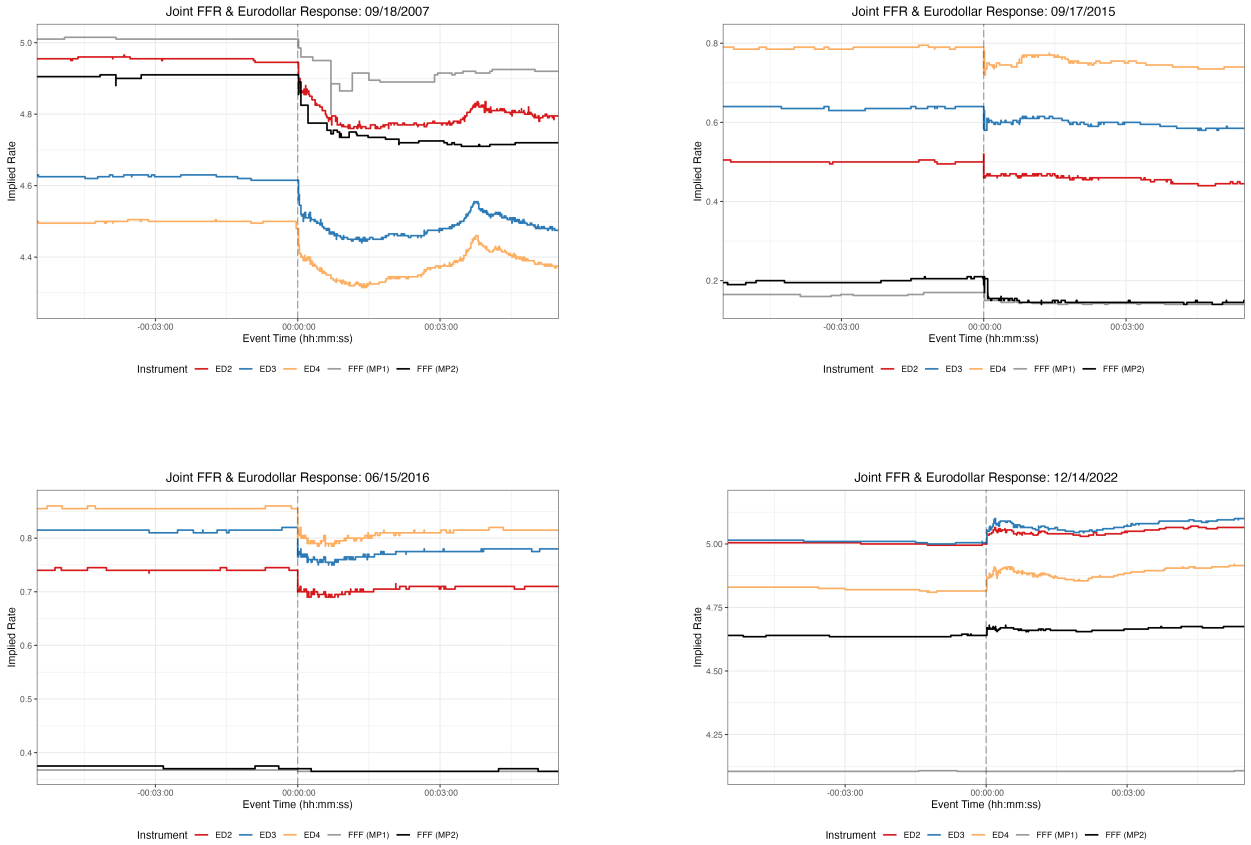


Figure 4: Four examples of Eurodollar reactions

These examples of price dynamics in narrow (10-minute) windows around announcements highlight several important features. First, even longer-maturity futures contracts can exhibit economically meaningful price changes almost instantaneously upon the announcement release. This rapid adjustment is consistent with mechanical algorithmic trading operating not only in the current-month contract but also across the term structure.

Second, the Eurodollar contracts tend to move almost in lockstep with one another, consistent with intensive calendar spread trading within Eurodollar contracts. Visually, the co-movement across Eurodollar maturities appears substantially stronger than the co-movement between Fed Funds Futures and Eurodollar contracts. For example, the correlation between MP2 and ED2 — contracts that span similar horizons — is noticeably weaker than the correlation between ED2 and ED4 — contracts that span more distinct horizons.[3] This pattern suggests a degree of market segmentation or imperfect cross-market integration between the Fed Funds and Eurodollar futures markets.

This distinction may be relevant for common approaches to estimating monetary policy shocks, which often rely on differences between near-term and longer-term futures responses to isolate forward guidance or information effects. Such strategies implicitly assume that differential responses across maturities reflect differential learning about future policy. However, the high-frequency correlation structure documented here suggests that microstructure features — possibly related to differences in liquidity and arbitrage linkages — may also contribute to the observed cross-contract differences.

3.1.3 Unconventional Policy Announcements

One way to further assess the interpretation advanced above is to examine FOMC announcements that contain information that is less straightforward to arbitrage mechanically. Announcements related to unconventional policy—such as large-scale asset purchases and forward guidance—provide a useful case study for this purpose.

Unlike announcements of the target policy rate, unconventional policy announcements primarily convey information through textual descriptions of asset purchase programs or qualitative guidance about the future path of rates. These details require language processing and economic interpretation to translate into implications for asset prices. Moreover, even conditional on understanding the announced policy, the mapping into equilibrium

[3]Figure A.11 in Appendix A.3 provides more direct evidence of this by plotting the correlation between the 10/20 meeting changes in MP2 vs. ED2 and ED2 vs. ED4, with the latter having a much stronger correlation.

yields is not mechanical. The academic literature reflects substantial uncertainty regarding both the transmission channels and quantitative effects of large-scale asset purchases (e.g., Kuttner (2018)), and similar ambiguity surrounds the strength and credibility of forward guidance. This uncertainty was particularly pronounced during QE1, which represented largely uncharted territory in the United States. For these reasons, it is less plausible that fully automated arbitrage strategies could convert such announcements into immediate price adjustments in the same way as conventional policy rate decisions.

To examine the temporal dynamics around these episodes, I plot implied interest rates from the futures contracts analyzed above for four QE1 announcements identified by Krishnamurthy et al. (2011): November 25 and December 16, 2008, and January 28 and March 18, 2009. I omit Chairman Bernanke's December 1 speech, as the timing of information release during a speech is inherently ambiguous. In addition, I examine the March 23, 2020 announcement.[4] For each episode, I plot a two-hour window surrounding the announcement to provide a broader view of the adjustment process. The results are shown in Figure 5.

The clearest cases are the standalone, unscheduled announcements on November 25, 2008, and March 23, 2020. In both instances, implied Eurodollar rates decline following the announcement, but the adjustment occurs gradually rather than instantaneously. Prices appear to take roughly 20 to 30 minutes to settle at their new levels. The remaining three announcements coincide with regularly scheduled FOMC meetings that also included policy rate decisions, making it more difficult to disentangle the sources of price movements. Nevertheless, by December 2008 the federal funds rate was at or near the effective lower bound and widely expected to remain there, suggesting limited uncertainty about the contemporaneous rate decision. Instead, these announcements primarily conveyed information about asset purchases and forward guidance. These price adjustments exhibit noticeable delays relative to the near-instantaneous responses documented for standard policy rate announcements in Figure 4. For example, the large adjustment on December 16, 2008 occurs only after approximately ten minutes. In other words, when the economic implications of an announcement

[4]I omit the March 15, 2020 announcement because it occurred on a Sunday when markets were closed.

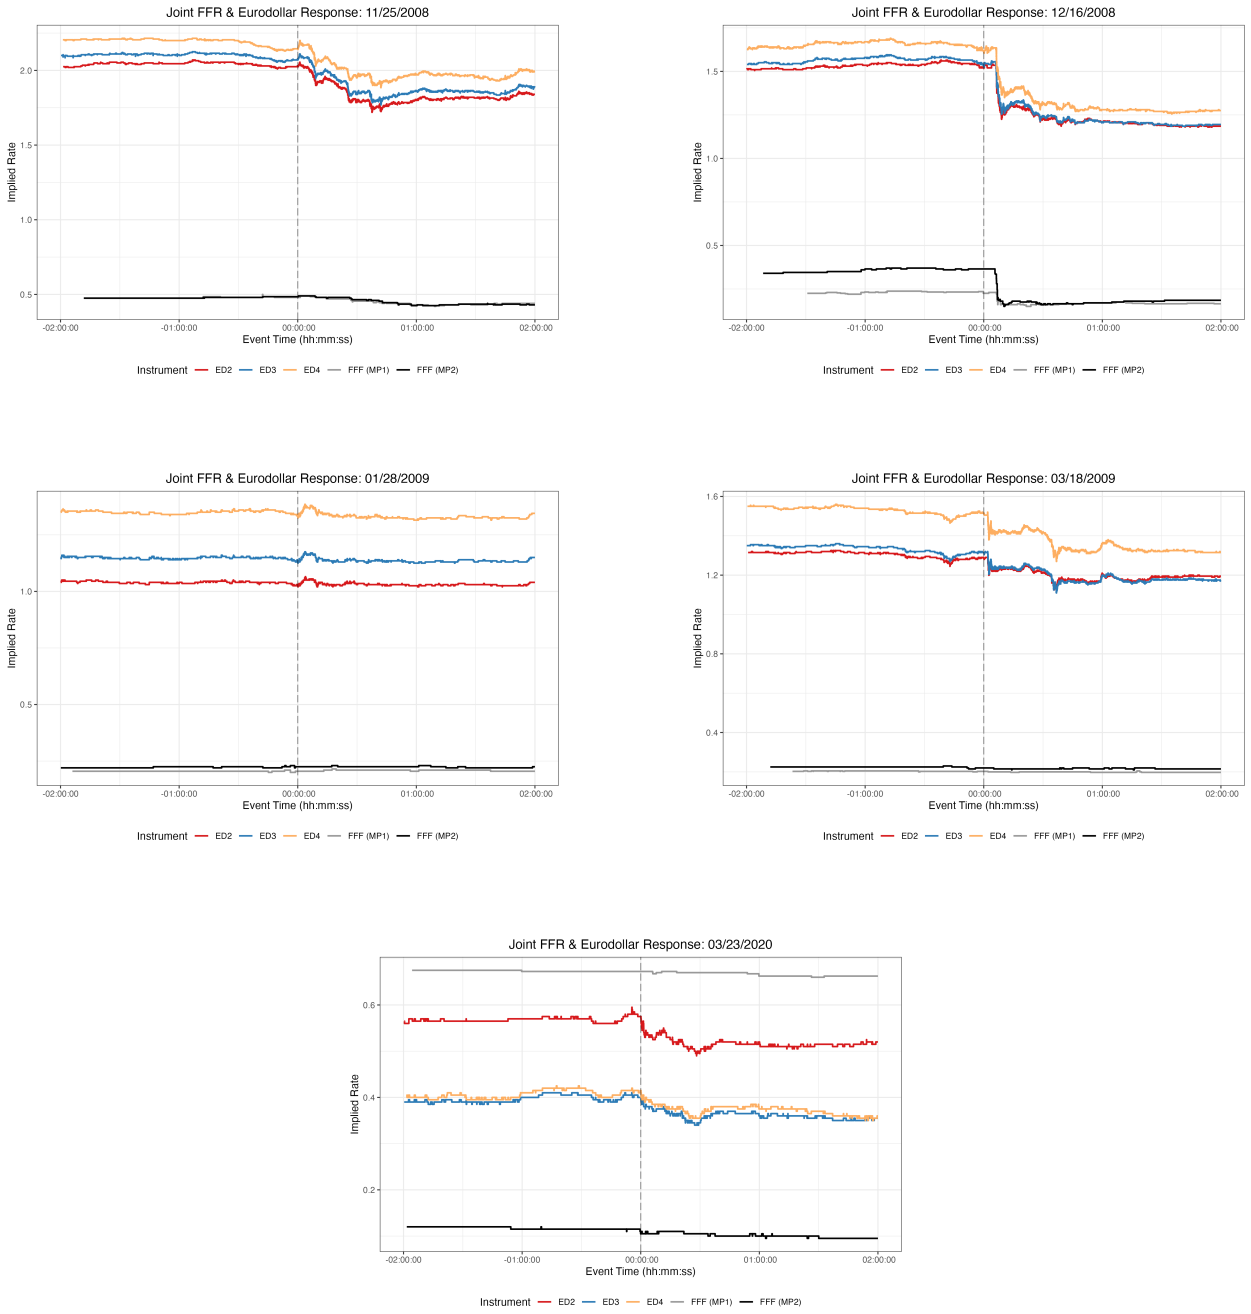


Figure 5: Future rate dynamics around unconventional policy announcements

are less mechanical and more interpretive, price discovery appears to operate on time scales that are plausibly consistent with human processing.[5]

Although the number of unconventional policy episodes is limited, the contrast in temporal dynamics is informative. When policy communication is text-heavy, interpretive, and not governed by a simple arbitrage relationship, price discovery appears to unfold over a substantially longer horizon. This stands in sharp contrast to the rapid adjustment observed for conventional policy rate announcements and is consistent with the view that when interpretation and economic reasoning are required, price adjustment occurs more gradually.

3.1.4 Speed of Price Discovery

The example FOMC meetings plotted above represent only a small subset of the more than 200 announcements in the sample. These examples raise two natural questions. First, how quickly does price discovery occur on average following an FOMC announcement? Second, has the speed of price discovery changed as financial markets have become increasingly automated?

One complication in addressing these questions is that many FOMC announcements produce little or no movement in futures prices. In such cases there is effectively no price discovery to measure. I therefore restrict the analysis to meetings where prices move meaningfully within the standard 10/20 window.

To characterize the typical pace of adjustment, I construct absorption measures for each meeting. Specifically, I calculate the time required for prices to realize 75% of the total change observed over the 10/20 window. Using these measures, I plot empirical cumulative distribution functions (ECDFs) of absorption times.

[5]The trading volume for these five announcements is plotted in Appendix A.2. Consistent with the price dynamics they generally do not show spikes in trading activity right at the announcement release. The November 25, 2008 and January 28th announcements show a gradual increase in volume. The December 16, 2008 and March 18, 2009 announcements do show sharp increases in trading volumes, but they both occur several minutes after the announcement, consistent with the timing of the price dynamics. Only the March 23, 2020 announcement shows a burst of trading right around the announcement. Looking closely at the price dynamics for the March 23, 2020 announcement, there is some volatility right at the announcement release, consistent with this volume increase, but the overall price change from this announcement unfolds more gradually.

I begin with the MP1 and MP2 contracts. Because these near-term federal funds futures contracts—particularly MP1—primarily reflect news about the current policy rate, I exclude meetings during the zero lower bound (ZLB) period, when policy announcements were more likely to convey information about future policy rather than the current target rate. To identify meetings with meaningful price responses, I require the absolute change in the implied interest rate over the 10/20 window to exceed 1 basis point. This slightly relaxed criteria from the 3bp used above helps increase the number of observations, but it remains relatively restrictive retaining fewer than half of the meetings in the full sample. To illustrate how the speed of price discovery has evolved over time, I plot the ECDFs separately across several sample eras. The results are shown in Figure 6.

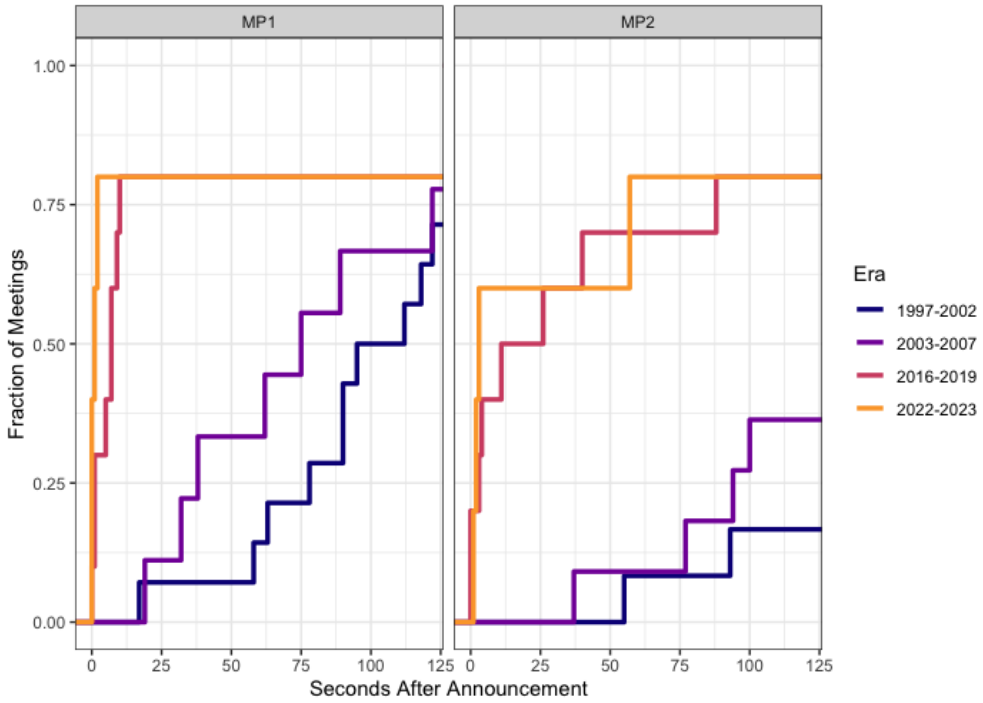


Figure 6: Fed Fund Futures: Price Discovery Speed Over Time

The absorption plots reveal several clear patterns. First, a large share of meetings exhibit extremely rapid price adjustment, with prices often reaching most of their total 10/20 change within seconds of the announcement. Second, the pace of price discovery has accelerated markedly over time, particularly in the post-2016 period. This timing coincides with

significant structural changes at the Chicago Mercantile Exchange, where these futures are traded. In 2015, the CME Group announced the closure of most open-outcry trading pits in favor of electronic trading (CME Group (2015b)) and introduced infrastructure upgrades aimed at reducing trading latency (CME Group (2015a)).

To provide a broader view of price discovery across the interest rate term structure, I repeat this analysis for the ED2–ED4 Eurodollar futures contracts. Because the price responses of these longer-maturity contracts are typically larger, I increase the threshold for "meaningful" announcements to 3 basis points. In addition, I plot meetings from the ZLB period as a separate ECDF. During the ZLB, market participants widely expected the federal funds rate to remain at zero, implying that most policy news likely concerned the future path of policy rather than the current target rate. As discussed earlier, information about future policy is more likely to require interpretation of FOMC communication and may therefore be incorporated more gradually into prices.

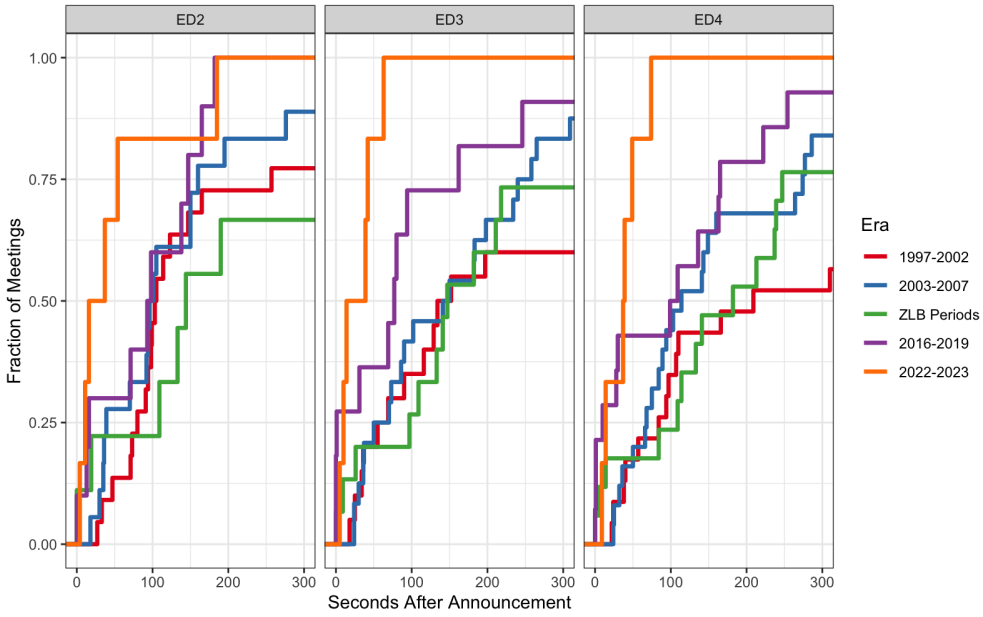


Figure 7: Eurodollar Futures: Price Discovery Speed Over Time

The Eurodollar results display patterns similar to those observed for federal funds futures. Price discovery has accelerated substantially over time, and a large share of announcements

exhibit substantial adjustment within the first minute. However, the ZLB meetings stand out for their slower dynamics, with absorption times often exceeding those observed even in the earliest years of the sample (1997–2002). This pattern is consistent with the examples of unconventional policy announcements shown earlier. When policy news primarily concerns the future path of monetary policy rather than the contemporaneous policy rate, price discovery tends to unfold more gradually, consistent with the need for market participants to interpret the informational content of the announcement.

3.2 Speed of Price Discovery and Shock Construction

The figures above document that prices in interest rate futures can respond extremely rapidly to FOMC announcements. Across both current-month federal funds futures and longer-maturity Eurodollar contracts, economically meaningful price adjustments often occur within seconds of the announcement release. At the same time, the examples of unconventional policy announcements illustrate that such rapid adjustment is not universal: when policy communication relies more heavily on textual guidance and interpretive content, price discovery can unfold over longer horizons. These contrasting patterns are broadly consistent with differences in the ease with which information can be processed and acted upon.

These results raise natural questions about the construction of monetary policy shocks, and what those shocks represent. In constructing shocks based on price changes in the 10/20 window commonly used in the literature, movements that occur within the first few seconds of the announcement are combined with adjustments that unfold over a longer horizon. As a result, the measured shock may reflect a mixture of different mechanisms. Some portion of the price change may arise from mechanical arbitrage relationships across closely related contracts that adjust almost immediately once the policy decision is revealed. Other movements may reflect a more deliberative process in which market participants interpret the broader policy communication and update expectations about the future path of policy.

The standard event window makes it difficult to distinguish between these channels. Yet this distinction matters for interpretation, as the literature often implicitly treats high-

frequency monetary policy shocks as capturing the market's reaction to newly revealed policy information. Examining how quickly the price movements underlying these shocks occur can therefore be informative about the nature of the adjustment process.

To explore this issue, I replicate the conventional approach in the literature to constructing monetary policy shocks, but vary the post-announcement window used to measure price changes. Specifically, I compute shocks using 15-second, 30-second, and 20-minute windows following the FOMC release. Comparing shocks constructed over these alternative horizons provides a systematic way to assess how much of the identified monetary policy shock is already embedded in prices within the earliest seconds after the announcement.

The logic of this exercise is straightforward. If the economically relevant component of price discovery occurs at speeds consistent with mechanical or algorithmic trading—particularly trading on the easily observable policy rate decision—then shocks constructed from very narrow windows should closely resemble those constructed using the conventional 20-minute window. In contrast, if market participants gradually incorporate more nuanced information from the announcement as they interpret its broader content, shocks measured over a 15-second window should differ meaningfully from those constructed over longer horizons.

3.2.1 Single Factor Shocks: Nakamura and Steinsson (2018)

In implementing this exercise, I follow the shock construction approach introduced by Nakamura and Steinsson (2018). This framework is well suited to the present analysis for two reasons. First, it relies exclusively on intraday futures prices, allowing for a precise examination of price dynamics at high frequencies.[6] Second, Nakamura and Steinsson (2018) interpret high-frequency monetary policy shocks as potentially reflecting a "Fed information effect," whereby FOMC announcements lead market participants to update beliefs not only about the path of monetary policy but also about broader economic fundamentals.

If such interpretive updating reflects human processing of the announcement text, it is unlikely to occur within a matter of seconds. Constructing shocks using ultra-narrow

[6]In contrast, several influential approaches in the literature rely on daily data (e.g., Bu et al. (2021)), which are ill suited for studying price discovery within minutes or seconds of an announcement.

windows therefore provides a way to restrict attention to price movements that are plausibly too rapid to reflect this type of learning. Comparing these shocks to those constructed using conventional 20-minute windows allows for a direct assessment of whether the empirical content of the shock series is consistent with the interpretation advanced in Nakamura and Steinsson (2018).

To identify unexpected changes in monetary policy, I follow the approach of Nakamura and Steinsson (2018) (replicated in Brennan et al. (2025)) and construct high-frequency shocks using tick-by-tick data for five interest rate instruments: current-month Fed Funds futures (MP1), next-month Fed Funds futures (MP2), and the second, third, and fourth Eurodollar futures contracts (ED2, ED3, ED4). Together, these instruments capture market expectations for the path of short-term interest rates from the immediate horizon out to approximately one year ahead. For the MP1 contract, I apply the standard scaling adjustment based on the day of the month of the announcement to account for the fact that Fed Funds futures settle on the average daily effective rate for the entire month.[7]

As noted above, the 10/20 timing window is the standard in the literature, which I adopt and extend to investigate the speed of price discovery. Specifically, I compute changes in implied rates over three alternative post-announcement windows: 15 seconds, 30 seconds, and 20 minutes. The pre-announcement baseline is fixed at 10 minutes prior to the release.[8]

The choice of a 15-second window is necessarily somewhat arbitrary—as are all event windows in this literature—but it is motivated by two considerations. First, there is modest uncertainty regarding the precise instant at which the announcement becomes observable in transaction data. Second, and more importantly, it is implausible that human traders could read, interpret, and act upon the language of an FOMC statement within 15 seconds of its release. Typical FOMC statements contain 300–400 words, implying a reading time

[7]Following the literature (e.g., Kuttner (2001)), I apply this scaling for announcements occurring before the final seven days of the month. For announcements in the last seven days, I instead use the unscaled next-month Fed Funds futures contract.

[8]A related comparison of narrower and wider windows is conducted in Gürkaynak et al. (2005), though their focus is on extending the window beyond the conventional 30-minute interval rather than tightening it to ultra-short horizons.

on the order of one to two minutes, not including interpretation or trade execution.[9] The maintained assumption is therefore that price movements occurring within 15 seconds are unlikely to reflect interpretive learning from the textual content of the statement.

Following Nakamura and Steinsson (2018), I extract monetary policy shocks by performing a principal component analysis on the vector of rate changes across the five instruments:

$$X_{t,\tau} = [\Delta MP1_{t,\tau}, \Delta MP2_{t,\tau}, \Delta ED2_{t,\tau}, \Delta ED3_{t,\tau}, \Delta ED4_{t,\tau}]',$$

where $\tau \in \{15s, 30s, 20m\}$ denotes the measurement window. I estimate the first principal component separately for each window and scale the resulting factor to have unit variance. The resulting series, MP_Factor_{τ}, represents the dominant common component of high-frequency interest rate surprises within each window.

Comparing MP_Factor_{15s} and MP_Factor_{20m} provides a direct assessment of whether including price movements that occur well after the announcement release materially alters the estimated policy shock. If interpretive learning plays an important role in the construction of the conventional 20-minute shock, then the factor estimated using a 15-second window should differ meaningfully from its 20-minute counterpart. To illustrate the relationship across windows, Figure 8 plots the shocks constructed using the 15-second and 20-minute windows.

The fitted linear relationship between the +15 second and +20 minute shocks has a coefficient of 0.87 and is significant at the 1%-level, which is plotted as a solid blue line in Figure 8. A 45 degree line is plotted in dashed red. These results highlight a strong correlation between monetary policy shocks constructed in a very narrow (15 second) versus the much wider (20 minute) window. These results highlight that the vast majority of the 10/20 shocks are realized almost instantaneously, which is suggestive of price effects driven by algorithmic trading, and strain the idea that these shocks reflect human learning from the language of the statements.

To assess the economic content of the shocks constructed from these different windows, I

[9] Average adult silent reading speeds are approximately 200–300 words per minute.

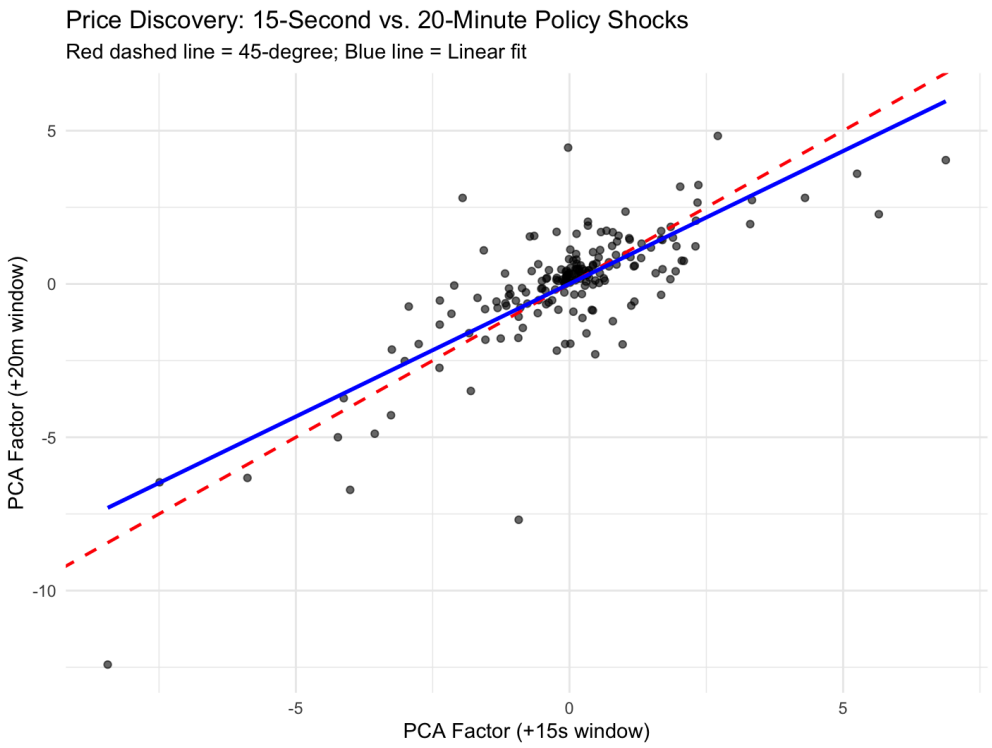


Figure 8: PCA Monetary Policy shocks from different windows

repeat the analysis on some primary outcome variables in Nakamura and Steinsson (2018). Specifically, I regress the three different shocks on the daily change in 3-month, 1-year, 5-year, and 10-year Treasury yields on the corresponding FOMC announcement days. These results are reported in Table I. It is important to note that the absolute magnitude of these regression coefficients is somewhat meaningless as they are the product of the scaling factor used in the PCA analysis. However, their significance and size across the different timing windows are more relevant.

	MP Factor 15sec	MP Factor 30sec	MP Factor 20min
3-Month	0.0047	0.0097*	0.0140**
	(0.0049)	(0.0046)	(0.0043)
1-Year	0.0167***	0.0159***	0.0177***
	(0.0026)	(0.0024)	(0.0022)
5-Year	0.0124**	0.0083*	0.0134***
	(0.0043)	(0.0041)	(0.0038)
10-Year	0.0060	0.0022	0.0089*
	(0.0044)	(0.0041)	(0.0039)

Notes: standard errors in parentheses. Significance levels: †p<0.1, *p<0.05, **p<0.01, ***p<0.001.

Table I: Regression Results: Impact of MP Factors on Yields

The results demonstrate a striking consistency in both the magnitude and significance of the coefficients across time horizons. For the 1-year and 5-year Treasury yields—the segments of the curve most sensitive to monetary policy—the impact of the 15-second shock is statistically indistinguishable from the 20-minute shock. For example, a one-unit increase in the 15-second factor is associated with a 1.67 basis point increase in the 1-year yield ($p < 0.001$), while the 20-minute factor yields a nearly identical coefficient of 1.77 basis point ($p < 0.001$). This stability suggests that the economic "content" of the monetary policy surprise is almost entirely captured in the immediate wake of the announcement. This finding casts further doubt on the hypothesis that these shocks primarily reflect a nuanced human interpretation of the FOMC statement text; if that were the case, one would expect the 20-minute window—which allows for the processing of qualitative language—to exhibit significantly different explanatory power than a window lasting only 15 seconds. It may be worth noting that these results are consistent with some other recent work challenging the importance of "Fed Information Effects" (Bauer and Swanson (2023)), but they do so for very different reasons.

These findings also point toward a deeper conceptual limitation in the single-factor Principal Component Analysis (PCA) used by Nakamura and Steinsson (2018), and others, to identify "Policy News." By construction, the first principal component seeks to maximize the explained variance across the full vector of instruments. Yet, as noted above, the front-

end instruments—specifically the current-month federal funds futures contract (MP1)—are arithmetically tied to the effective federal funds rate for the current month and therefore primarily reflect the realized policy rate decision rather than expectations about more distant policy outcomes.

When PCA is used to extract a single common factor across both "target" instruments such as MP1 and forward-looking instruments such as Eurodollar futures, the resulting factor is naturally anchored to the policy rate surprise that these contracts share. Because MP1 cannot independently respond to news about the distant path of policy, the first principal component will tend to place substantial weight on the immediate policy rate surprise.

This feature of the methodology also helps explain the similarity in shocks constructed using very narrow and conventional event windows. If the dominant source of common variation across instruments is the surprise in the current policy rate, then much of the variance captured by the single-factor shock will be realized almost immediately following the announcement. More broadly, this highlights a conceptual limitation of the single-factor PCA approach. When the instrument set includes contracts such as MP1 that primarily reflect the current policy rate, the resulting factor will naturally be dominated by policy rate surprises that all instruments share. Under these conditions, interpreting the single extracted factor as capturing "Policy News" or information about the future path of policy becomes problematic, as the methodology itself mechanically anchors the factor to the immediate policy decision.

3.2.2 Two-Factor Approach: Gürkaynak et al. (2005)

While the single-factor approach of Nakamura and Steinsson (2018) provides a parsimonious measure of the "Policy News" bundle, the results in the previous section—specifically the extreme stability of coefficients between the 15-second and 20-minute windows—suggest a potential methodological conflation. Because the first principal component is mathematically forced to explain the maximum shared variance across all instruments, it becomes heavily anchored by the current-month Fed Funds future (MP1). As noted, MP1 is arithmetically

tied to the immediate policy rate and is, by construction, invariant to expectations regarding the medium-term path. Consequently, if the "Policy News" is dominated by the headline rate decision, a single-factor model may mechanically drown out any distinct information contained in the FOMC statement, even when that information is available in the wider 20-minute window.

To address this potential conflation and allow for a multidimensional response to Fed communications, I implement the two-factor identification strategy proposed by Gürkaynak et al. (2005).[10] This approach explicitly recognizes that the Federal Reserve possesses two distinct tools: the "Target" (the immediate policy action) and the "Path" (forward guidance regarding future policy). By extracting the first two principal components and performing a structural rotation, their approach isolates a "Path Factor" that is restricted to have no impact on the current-month Fed Funds rate. This decomposition provides a more rigorous test of the price discovery process: if the Path truly represents the market's deliberative learning from the qualitative language of the FOMC statement, we should observe a distinct decoupling between the Path factor measured at 15 seconds and the version realized after 20 minutes of human processing.

The decomposition of monetary policy surprises into Target and Path factors allows for a more granular examination of how different types of information are absorbed by the market. If the Path factor represents the market's interpretation of the FOMC statement's forward-looking language, we should observe a significant evolution in this factor between the 15-second and the 20-minute windows. Figure 9 presents a scatter plot of the GSS Path factor constructed at the 15-second horizon against the 20-minute horizon. The solid blue line represents the linear fit, while the dashed red line represents the 45-degree identity line. The relationship is positive and highly significant, with a regression coefficient of 0.66 ($p < 0.001$) and an R^2 of 0.38.The fact that nearly 40% of the variation in the 20-minute Path factor is already present within 15 seconds of the announcement poses a significant

[10]There are more recent iterations of this general approach (e.g. Bauer and Swanson (2023)) but I focus on this original framework as much of the later work in the literature builds from this foundation, and the analysis in Gürkaynak et al. (2005) is more parsimonious than later work, making it easier to disentangle the factors driving the results.

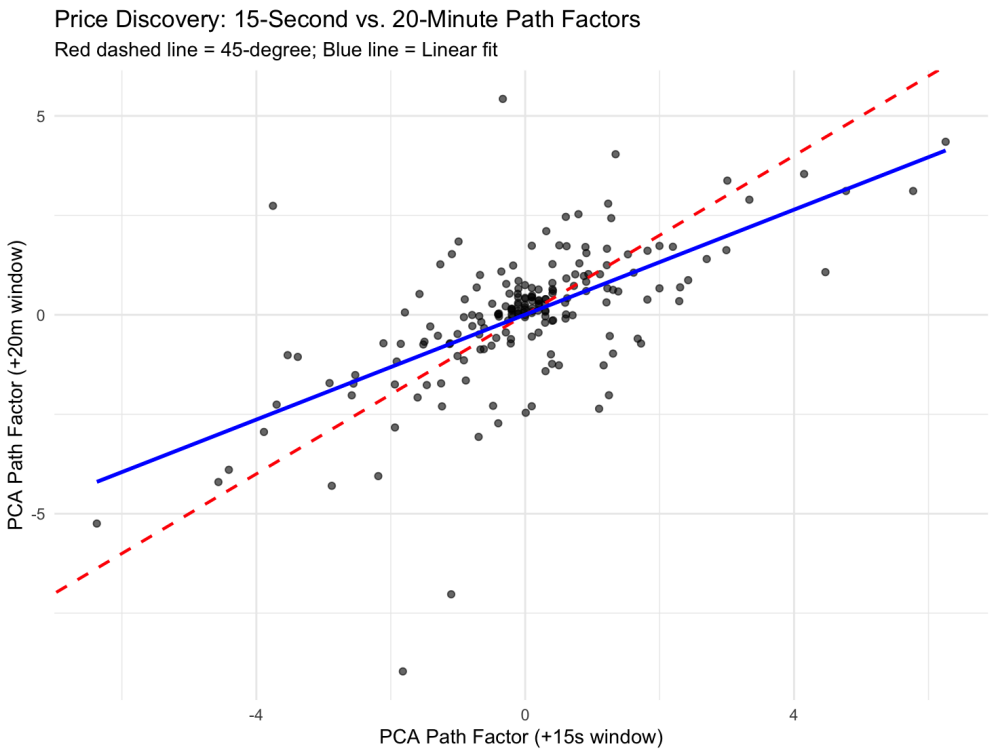


Figure 9: PCA Path Factor from different windows

challenge to the standard interpretation of forward guidance shocks. Because it is unlikely that human agents to have parsed, discussed, and traded upon the qualitative nuances of the FOMC statement within 15 seconds, these results suggest that a substantial portion of the "Path" is realized via algorithmic responses to the headline figures. This "Instant Path" likely reflects the automated adjustment of the Eurodollar strip—where liquidity is often higher and price increments more granular than in the Fed Funds futures market—rather than a nuanced learning process regarding the Fed's future intentions.

	GSS Path 15 sec	GSS Path 30 sec	GSS Path 20 min
3-Month	0.0022	0.0080	0.0121*
	(0.0054)	(0.005)	(0.005)
1-Year	0.0153***	0.0155***	0.018***
	(0.0029)	(0.0028)	(0.0026)
5-Year	0.0148**	0.011*	0.0189***
	(0.0046)	(0.0046)	(0.0042)
10-Year	0.0090†	0.0048	0.014**
	(0.0048)	(0.0047)	(0.0043)

Notes: standard errors in parentheses. Significance levels: †p<0.1, *p<0.05, **p<0.01, ***p<0.001.

Table II: The Impact of GSS Path Factors on Treasury Yields

The regression results in Table II provide a nuanced view of how the Path factor impacts the Treasury curve across maturities and event windows. Consistent with the standard interpretation of a Path shock, the effects are concentrated beyond the very front end of the curve: the 3-month bill shows little response in the narrowest windows, while the 1-year, 5-year, and 10-year yields respond significantly.

Crucially, however, the "Instant Path" identified within 15 seconds of the announcement already carries substantial economic content. For the 1-year Treasury note, the coefficient for the 15-second Path factor (0.0153) is close to the 20-minute coefficient (0.018), and both are highly significant ($p < 0.001$). Similar patterns hold at longer maturities: at the 5-year horizon, the 15-second coefficient (0.0148) is economically large relative to the 20-minute estimate (0.0189), and at the 10-year horizon the 15-second response (0.0090) accounts for a sizable share of the 20-minute effect (0.014). Although the 20-minute coefficients are

uniformly larger, a substantial portion of the total response is already realized in the first few seconds following the announcement.

3.3 Monetary Policy Shocks and Limits to Arbitrage

The evidence presented above suggests that the content of monetary policy shocks identified using high-frequency approaches may, at least in part, reflect mechanical arbitrage relationships across closely related interest rate contracts. A large literature in finance documents that such arbitrage relationships can weaken during periods of financial market stress, particularly when intermediary balance sheets are constrained (e.g., Du et al. (2018)). This raises the possibility that limits to arbitrage may influence the construction of monetary policy shocks. To explore this possibility, I examine whether commonly used measures of financial market stress are related to the magnitude of both policy rate shocks (Target) and shocks affecting later-expiring futures (Path).[11]

3.3.1 Limits to arbitrage and Target shocks

The justification for using forward contracts to construct monetary policy shocks is fundamentally based on an no-arbitrage argument: the forward price (implied interest rate) immediately prior to an FOMC announcement should reflect the market's expectation of the forthcoming policy decision. If this were not the case, market participants could in principle exploit the mispricing by buying or selling the relevant futures contracts.

In practice, however, arbitraging such discrepancies ahead of an FOMC announcement requires bearing risk. Market participants must commit capital while facing uncertainty about the policy decision itself, as well as potential counterparty, funding, or liquidity risks. During periods of financial stress, when balance sheet capacity is limited, arbitraging small pricing discrepancies may not be an attractive use of scarce regulatory capital. Once the FOMC decision is announced, uncertainty about the policy outcome is resolved and the risk associated with these positions largely disappears.

[11] Bauer and Swanson (2023) also explore how monetary policy shocks correlate with pre-announcement economic and financial market data, but they focus on different financial market data than I do here.

If limits to arbitrage are relevant in this setting, then the magnitude of measured Target shocks may vary systematically with conditions in financial markets. Periods of elevated financial stress may be associated with larger measured shocks, reflecting a greater degree of pre-announcement mispricing that is resolved when the policy decision becomes public.

To explore this question, I use three different measures of financial market stress. To capture different dimensions of financial market stress, I use three widely studied indicators: the ICE BofA MOVE Index, the illiquidity measure proposed by Hu et al. (2013) (often referred to as the "noise as information" measure), and the TED spread. These indicators capture complementary aspects of stress in financial markets, including uncertainty about interest rates, market liquidity, and funding conditions in the banking sector.

The MOVE index is often described as the bond-market analogue to the equity volatility index CBOE Volatility Index (VIX). It measures the implied volatility embedded in options on U.S. Treasury securities across several maturities. Higher values of the MOVE index indicate elevated uncertainty about the future path of interest rates and Treasury yields. Periods of heightened interest rate volatility may limit the willingness of arbitrageurs to take positions that exploit small deviations in interest rate derivatives pricing, particularly when such trades require balance sheet capacity or expose traders to mark-to-market risk.

The "noise as information" measure developed by Hu et al. (2013) is designed to capture mispricing and illiquidity in the U.S. Treasury market. The measure is constructed by comparing observed Treasury yields to those implied by a fitted yield curve model. In liquid and well-functioning markets, arbitrage relationships across Treasury securities should keep these deviations small. Larger deviations therefore reflect either limits to arbitrage or frictions that prevent traders from enforcing the no-arbitrage relationships implied by standard term structure models. Elevated values of this measure are typically interpreted as evidence of reduced market liquidity or constrained arbitrage capacity in Treasury markets.

Finally, the TED spread measures the difference between the three-month interbank lending rate and the yield on three-month U.S. Treasury bills. Historically, it has been interpreted as a proxy for perceived credit risk and funding stress in the banking sector.

A wider TED spread indicates that banks face higher costs of unsecured funding relative to the risk-free rate. Because many arbitrage strategies in interest rate derivatives rely on intermediary balance sheets and short-term funding, elevated funding stress may reduce the ability or willingness of financial intermediaries to exploit mispricing in futures and forward markets.

Together, these measures provide a broad view of financial market conditions. The MOVE index captures uncertainty about interest rates, the noise measure captures liquidity and pricing dislocations in Treasury markets, and the TED spread reflects stress in short-term funding markets. If limits to arbitrage play an important role in the construction of high-frequency monetary policy shocks, one might expect the relationship between futures prices and the underlying policy expectations to weaken during periods when these indicators signal heightened financial market stress.

To examine whether financial market frictions are associated with the magnitude of identified monetary policy shocks, I regress the absolute value of the GSS Target factor on several measures of financial market stress. Because monetary policy surprises can be either positive or negative, using the absolute value of the shock focuses on the magnitude of the announcement response rather than its direction. If limits to arbitrage weaken the link between futures prices and underlying policy expectations prior to FOMC announcements, one might expect larger measured shocks during periods of heightened financial market stress.

I estimate a series of regressions relating the magnitude of the Target factor to three indicators of financial stress: the ICE BofA MOVE Index, the Hu et al. (2013) Noise Measure, and the TED spread. To control for any possible endogenous response of these factors to FOMC decisions, I use their value on the day before each FOMC announcement. Each measure is first considered individually. Because these indicators capture related but distinct aspects of financial market stress and are moderately correlated with one another, I also estimate a specification that includes all three simultaneously. The results are reported in the top panel of Table III.

The results in Panel A indicate that financial market conditions explain a substantial

Panel A: GSS Target	(1)	(2)	(3)	(4)
Dep. Var: $\lvert Target_{20min} \rvert$	*TED Only*	*MOVE Only*	*Noise Only*	*Full Model*
TEDRATE_{t-1}	1.262***			1.324***
	(0.169)			(0.228)
MOVE_{t-1}		0.013***		0.010**
		(0.003)		(0.004)
Noise_{t-1}			0.128***	-0.133*
			(0.034)	(0.052)
R-squared	0.270	0.153	0.086	0.310
Panel B: GSS Path	(5)	(6)	(7)	(8)
Dep. Var: $\lvert Path_{20min} \rvert$	*TED Only*	*MOVE Only*	*Noise Only*	*Full Model*
TEDRATE_{t-1}	0.603**			-0.068
	(0.183)			(0.233)
MOVE_{t-1}		0.015***		0.016***
		(0.002)		(0.004)
Noise_{t-1}			0.143***	-0.010
			(0.032)	(0.053)
R-squared	0.067	0.214	0.118	0.215

Notes: Standard errors in parentheses. *p<0.05, **p<0.01, ***p<0.001.

Table III: Institutional Frictions and Monetary Policy Shocks (GSS Factors)

fraction of the variation in the magnitude of policy rate surprises. When considered individually, the TED spread alone explains roughly 27 percent of the variation in the absolute value of the Target factor. The MOVE index also has meaningful explanatory power, accounting for approximately 15 percent of the variation, while the Treasury market noise measure explains roughly 9 percent.

When all three measures are included simultaneously, the R^2 increases modestly to 0.31. This suggests that a significant portion of the variation in Target shocks is related to broader financial market conditions, with funding market stress appearing to play a particularly important role. The relatively small increase in explanatory power when adding the MOVE and noise measures to the TED specification indicates that these indicators capture overlapping dimensions of financial stress.

The strong explanatory power of the TED spread is particularly notable. The TED spread measures the difference between unsecured interbank lending rates and short-term Treasury yields and is commonly interpreted as an indicator of stress in short-term funding markets. Elevated values of the TED spread reflect periods when banks face higher costs of unsecured funding and tighter balance sheet constraints. Because arbitrage strategies in interest rate futures markets are typically implemented by financial intermediaries and rely on short-term funding and balance sheet capacity, stress in funding markets may limit the willingness or ability of arbitrageurs to enforce pricing relationships between futures contracts and underlying policy expectations prior to FOMC announcements. In such environments, deviations between futures prices and policy expectations may persist until the announcement resolves uncertainty, resulting in larger measured Target shocks.

Taken together, these results suggest that a nontrivial share of the measured variation in policy rate surprises is associated with conditions in financial markets, consistent with the idea that limits to arbitrage or intermediary balance sheet constraints affect the extent to which futures prices reflect policy expectations prior to FOMC announcements.

3.3.2 Limits to arbitrage and Path shocks

The standard PCA to constructing Path shocks isolates the component of price movements in later-expiring (Eurodollar) futures contracts that is orthogonal to movements in the near-expiring federal funds futures contract. In practice, the Path factor reflects differential movements along the interest rate futures curve, capturing changes in market expectations about the future path of policy rates beyond the current FOMC meeting.

As with the Target factor, the interpretation of Path shocks implicitly relies on arbitrage relationships across futures contracts of different maturities. In well-functioning markets, movements in the front contract should transmit smoothly to longer-dated contracts through arbitrage trading along the futures curve. However, enforcing these relationships requires market participants to intermediate duration risk and commit balance sheet capacity across multiple maturities.

In contrast to the Target factor, which primarily reflects pricing in the contract most closely tied to the upcoming policy decision, the Path factor depends on pricing relationships across multiple maturities. In particular, an arbitrageur attempting to align the response of longer-dated contracts (such as ED4) with movements in the front contract (MP1) must take positions that expose them to greater interest rate risk and mark-to-market volatility. When financial intermediaries face tighter funding conditions or balance sheet constraints, their willingness to take on these positions may be limited. As a result, the pass-through of policy shocks across the futures curve may weaken, leading to larger differential movements between near- and longer-dated contracts. If the Path factor reflects these differential responses, its magnitude may therefore be sensitive to the prevailing risk environment and the associated shadow cost of intermediary balance sheet capacity.

The results for the Path factor are presented in Panel B of Table III. As with the Target factor, each measure of financial stress is positively and significantly related to the magnitude of Path shocks when considered individually. However, the relative importance of the different indicators differs noticeably from the pattern observed for policy rate surprises.

In particular, the ICE BofA MOVE Index appears to play the dominant role. When

considered on its own, the MOVE index explains roughly 21 percent of the variation in the magnitude of Path shocks, substantially more than either the TED spread or the Treasury noise measure. Moreover, when all three measures are included simultaneously, the MOVE index remains highly statistically significant and retains essentially all of its explanatory power, while the coefficients on the TED spread and the Hu et al. (2013) Noise Measure become small and statistically insignificant.

This pattern is consistent with the mechanism discussed above. Interpreting movements in longer-dated futures contracts relative to the front contract requires arbitrage across the term structure of interest rate futures. Implementing such trades exposes intermediaries to duration risk and mark-to-market volatility. As a result, the willingness of market participants to enforce pricing relationships across maturities may depend more heavily on the level of interest rate uncertainty than on conditions in short-term funding markets. Periods of elevated interest rate volatility—captured by higher values of the MOVE index—may therefore weaken the transmission of policy news across the futures curve, leading to larger differential movements between near- and longer-dated contracts and, consequently, larger measured Path shocks.

Taken together, the results in Panels A and B suggest that different dimensions of financial market stress matter for different components of high-frequency monetary policy shocks. Funding market stress appears most closely related to the magnitude of policy rate surprises, while interest rate volatility plays a more prominent role in explaining variation in shocks to the expected future path of policy. This distinction is consistent with the idea that enforcing pricing relationships in the front contract primarily relies on short-term funding and intermediary balance sheets, whereas arbitrage along the futures curve requires intermediating greater interest rate risk.

4 Discussion

This paper examines the timing of price discovery in interest rate futures following FOMC announcements. Using tick-by-tick data for Federal Funds and Eurodollar futures, I doc-

ument that economically meaningful price adjustments frequently occur within seconds of the announcement release. Shocks constructed using an ultra-narrow 15-second window are therefore highly correlated with shocks constructed using the conventional 10-minute-before to 20-minute-after window that is standard in the literature, and the two measures exhibit very similar relationships with Treasury yields across maturities.

These findings raise questions about the common interpretation of high-frequency monetary policy shocks. Much of the literature assumes that price adjustments within the standard event window partly reflect market participants interpreting the qualitative information embedded in FOMC communication—such as forward guidance, signals about the future policy path, or information about the central bank's economic outlook. However, the speed of the observed price responses suggests that a substantial portion of the adjustment occurs too quickly to plausibly reflect deliberative human interpretation of the announcement text. Instead, the dynamics are more consistent with algorithmic trading and rapid arbitrage across closely related futures contracts triggered by the immediately observable components of the announcement.

The contrast with unconventional policy episodes is informative in this regard. Announcements related to asset purchases or forward guidance—where the economic implications must be inferred from textual descriptions rather than a discrete policy rate decision—exhibit noticeably slower price discovery. In these cases, futures prices often take several minutes to converge to their new equilibrium levels, consistent with the idea that interpretation and belief updating require time.

Taken together, the results suggest that high-frequency monetary policy shocks measured in conventional event windows are, at least in part, the product of mechanical arbitrage across interest rate futures contracts. This observation motivates a second set of findings: if arbitrage is operative, then limits to arbitrage should matter for shock magnitudes. The paper confirms this directly. Measures of financial market stress explain a substantial share of the cross-meeting variation in both target-rate and path-factor shock magnitudes. The relevant frictions differ across shock types in a theoretically coherent way: funding market

stress is the primary driver of target-rate shock magnitude, consistent with the reliance of front-contract arbitrage on short-term funding and intermediary balance sheets, while interest rate uncertainty drives path-factor shock magnitudes, consistent with the greater duration exposure required to arbitrage along the futures curve.

Recognizing the role of mechanical arbitrage—and its limits—in the construction of high-frequency monetary policy shocks may help clarify several puzzling features of the literature. Shocks identified in conventional event windows may be larger during periods of financial stress not because policy was more surprising, but because pre-announcement mispricing was more severe or post-announcement arbitrage was less efficient. This raises the possibility that the measured magnitude of high-frequency monetary policy shocks is partly an artifact of market conditions rather than a clean signal of the information content of the announcement. Disentangling these forces is an important direction for future research.

References

Acosta, M. (2022): "The perceived causes of monetary policy surprises," *Published manuscript*.

Acosta, M., C. M. Brennan, and M. M. Jacobson (2024): "Constructing high-frequency monetary policy surprises from SOFR futures," *Economics Letters*, 242, 111873.

Bauer, M. D. and E. T. Swanson (2023): "A reassessment of monetary policy surprises and high-frequency identification," *NBER Macroeconomics Annual*, 37, 87–155.

Brennan, C. M., M. M. Jacobson, C. Matthes, and T. B. Walker (2025): "Monetary Policy Shocks: Data or Methods?" *The BE Journal of Macroeconomics*, 25, 595–659.

Bu, C., J. Rogers, and W. Wu (2021): "A unified measure of Fed monetary policy shocks," *Journal of Monetary Economics*, 118, 331–349.

CME Group (2015a): "CME Globex Notices: May 4, 2015," CME Group Electronic Trading Advisories, accessed: 2024-05-22.

——— (2015b): "CME Group to Close Most Open Outcry Futures Trading in Chicago and New York by July; Most Options Markets to Remain Open," CME Group Investor Relations, accessed: 2024-05-22.

Du, W., A. Tepper, and A. Verdelhan (2018): "Deviations from covered interest rate parity," *The Journal of Finance*, 73, 915–957.

Gürkaynak, R. S., B. Sack, and E. T. Swansonc (2005): "Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements," *International Journal of Central Banking*.

Hu, G. X., J. Pan, and J. Wang (2013): "Noise as information for illiquidity," *The Journal of Finance*, 68, 2341–2382.

Krishnamurthy, A., A. Vissing-Jorgensen, S. Gilchrist, and T. Philippon (2011): "The Effects of Quantitative Easing on Interest Rates: Channels and Implications for Policy/Comments and Discussion," *Brookings Papers on Economic Activity*, 215.

Kuttner, K. N. (2001): "Monetary policy surprises and interest rates: Evidence from the Fed funds futures market," *Journal of monetary economics*, 47, 523–544.

——— (2018): "Outside the box: Unconventional monetary policy in the great recession and beyond," *Journal of Economic Perspectives*, 32, 121–146.

Miller, R. S. and G. Shorter (2016): *High frequency trading: Overview of recent developments*, vol. 4, Congressional Research Service Washington, DC.

Nakamura, E. and J. Steinsson (2018): "High-frequency identification of monetary non-neutrality: the information effect," *The Quarterly Journal of Economics*, 133, 1283–1330.

Paul, P. (2020): "The time-varying effect of monetary policy on asset prices," *Review of Economics and Statistics*, 102, 690–704.

A Additional Results

A.1 Volume dynamics

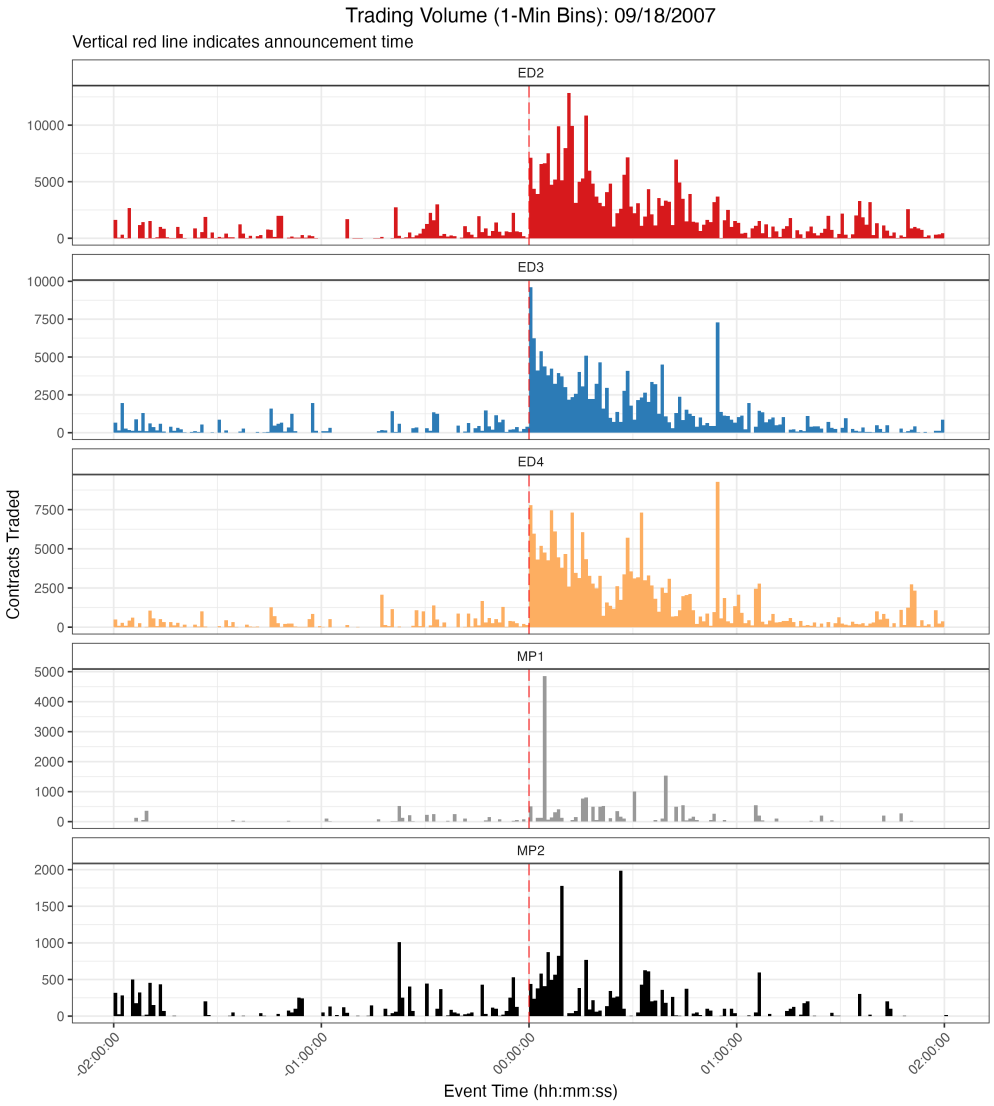


Figure A.1: Trading volume aggregated to 1-minute bins: 09/18/2007

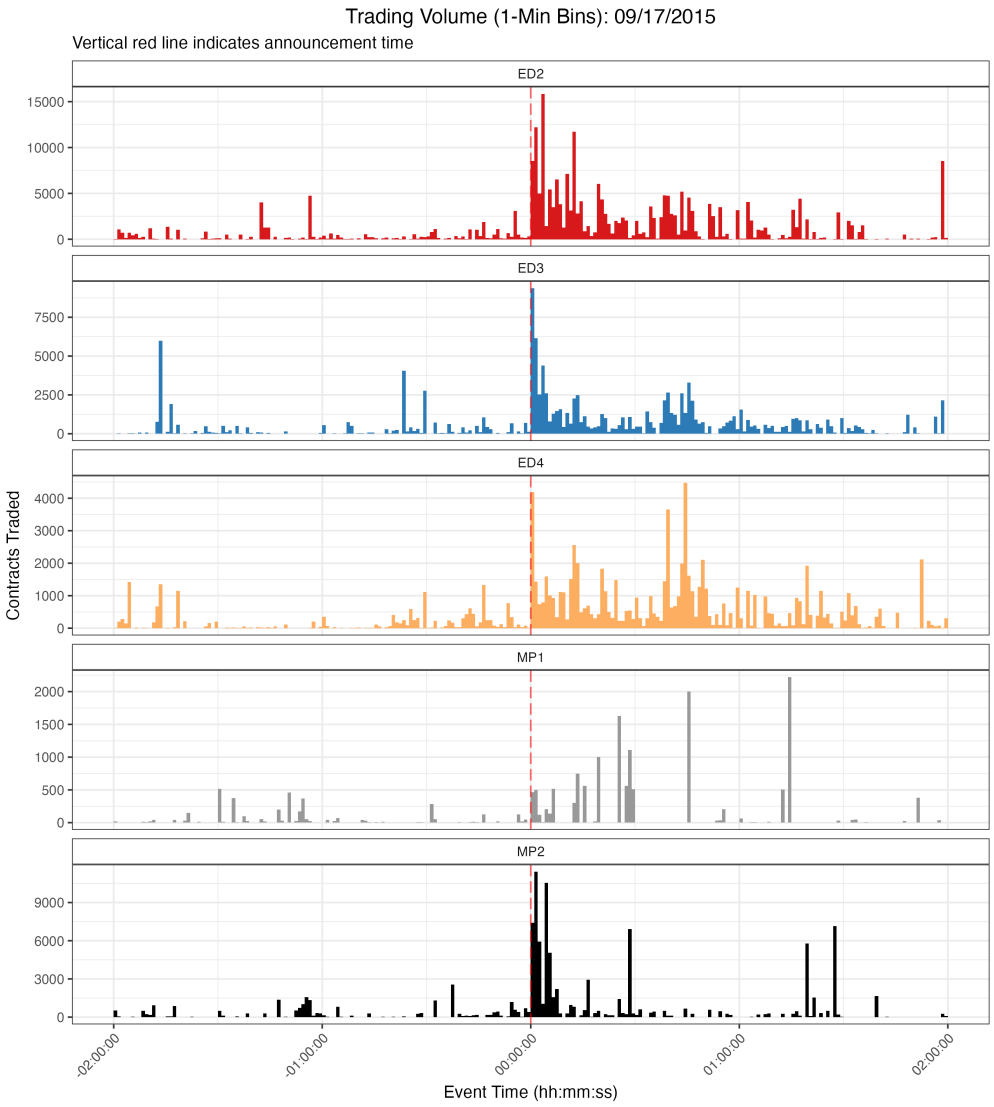


Figure A.2: Trading volume aggregated to 1-minute bins: 09/17/2015

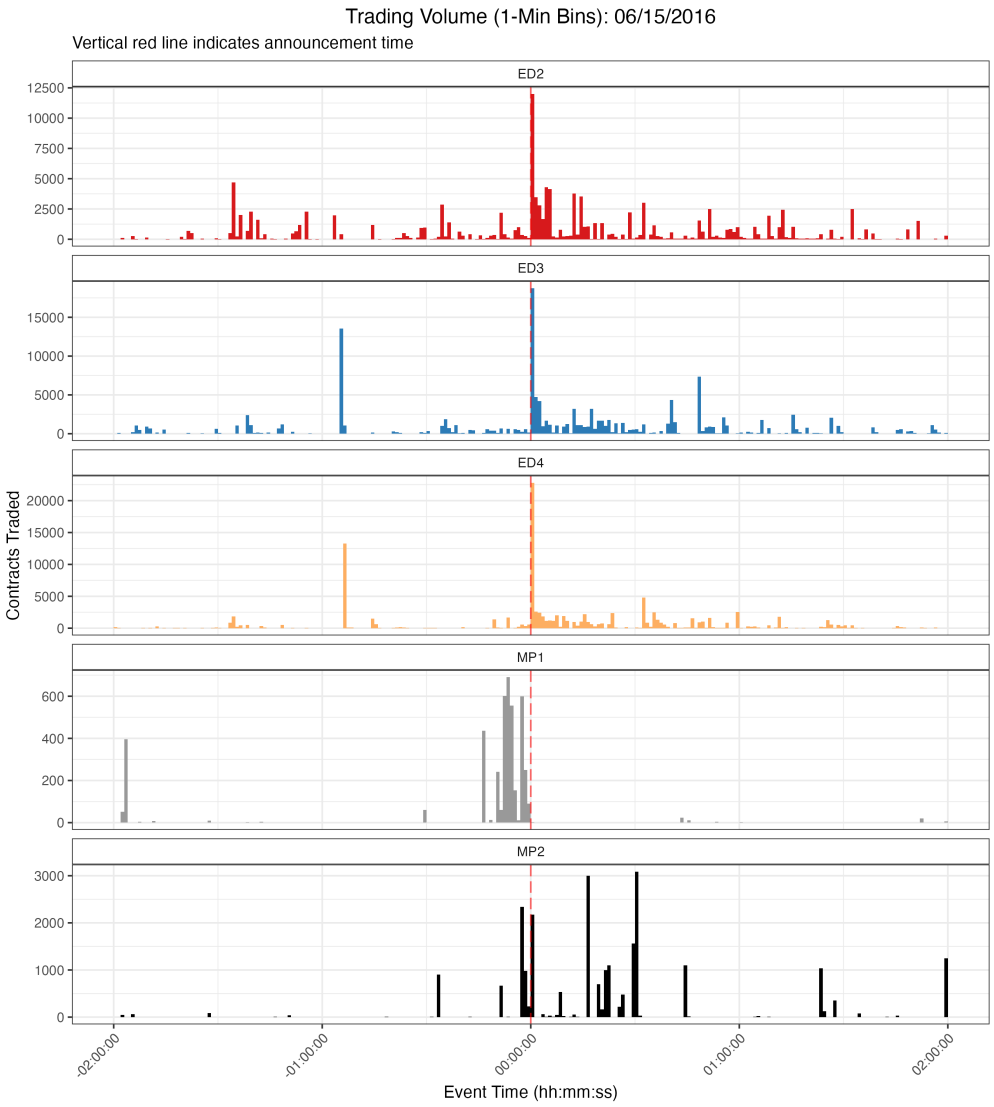


Figure A.3: Trading volume aggregated to 1-minute bins: 05/15/2016

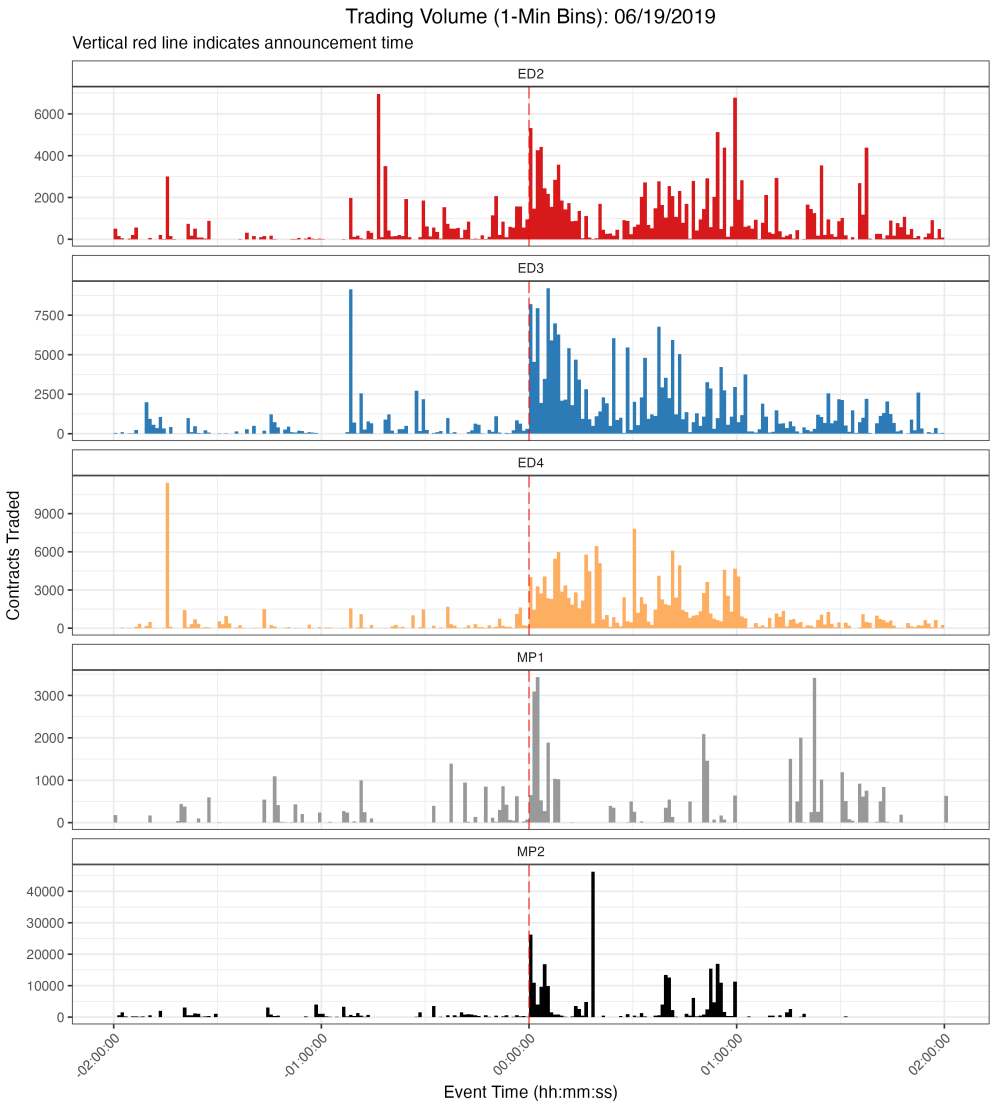


Figure A.4: Trading volume aggregated to 1-minute bins: 06/19/2019

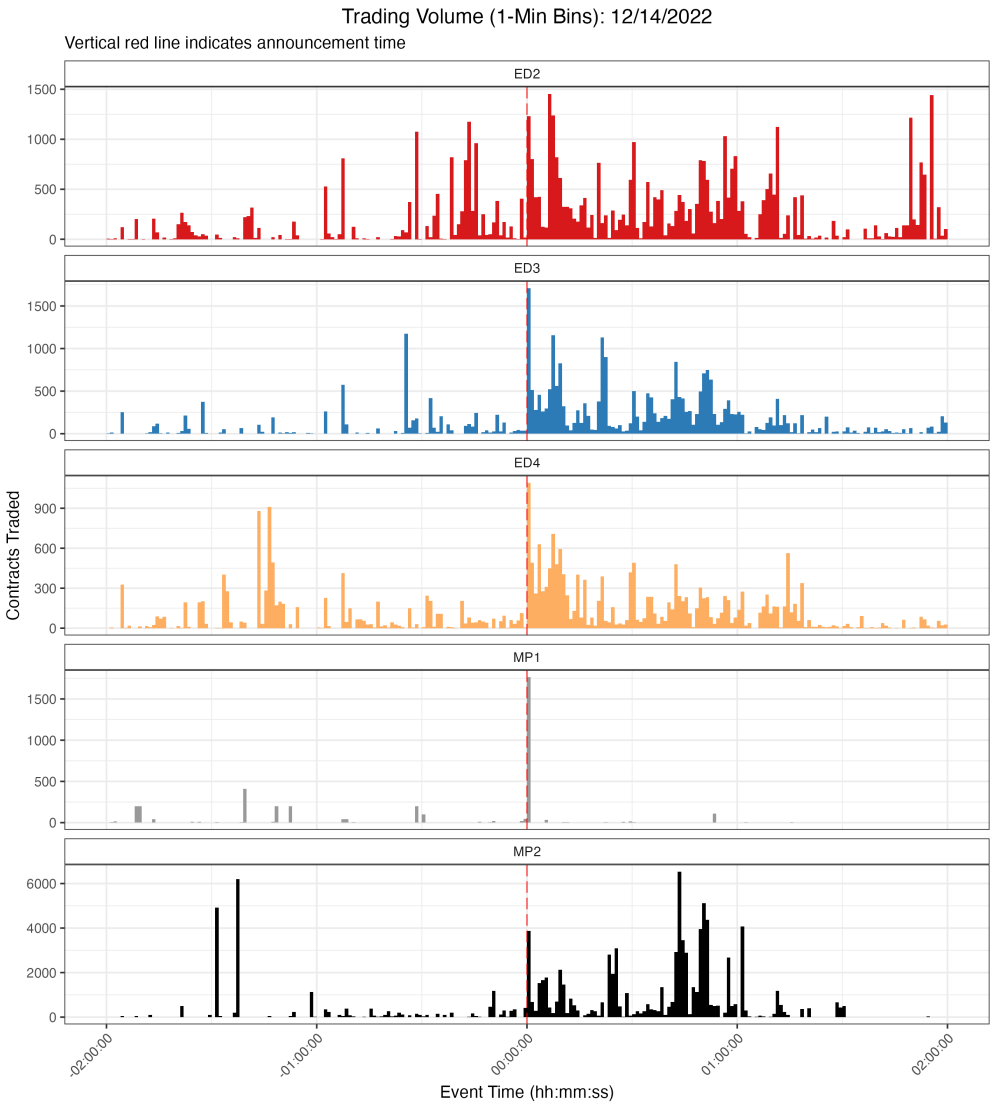


Figure A.5: Trading volume aggregated to 1-minute bins: 14/12/2022

A.2 Volume dynamics around important QE announcements

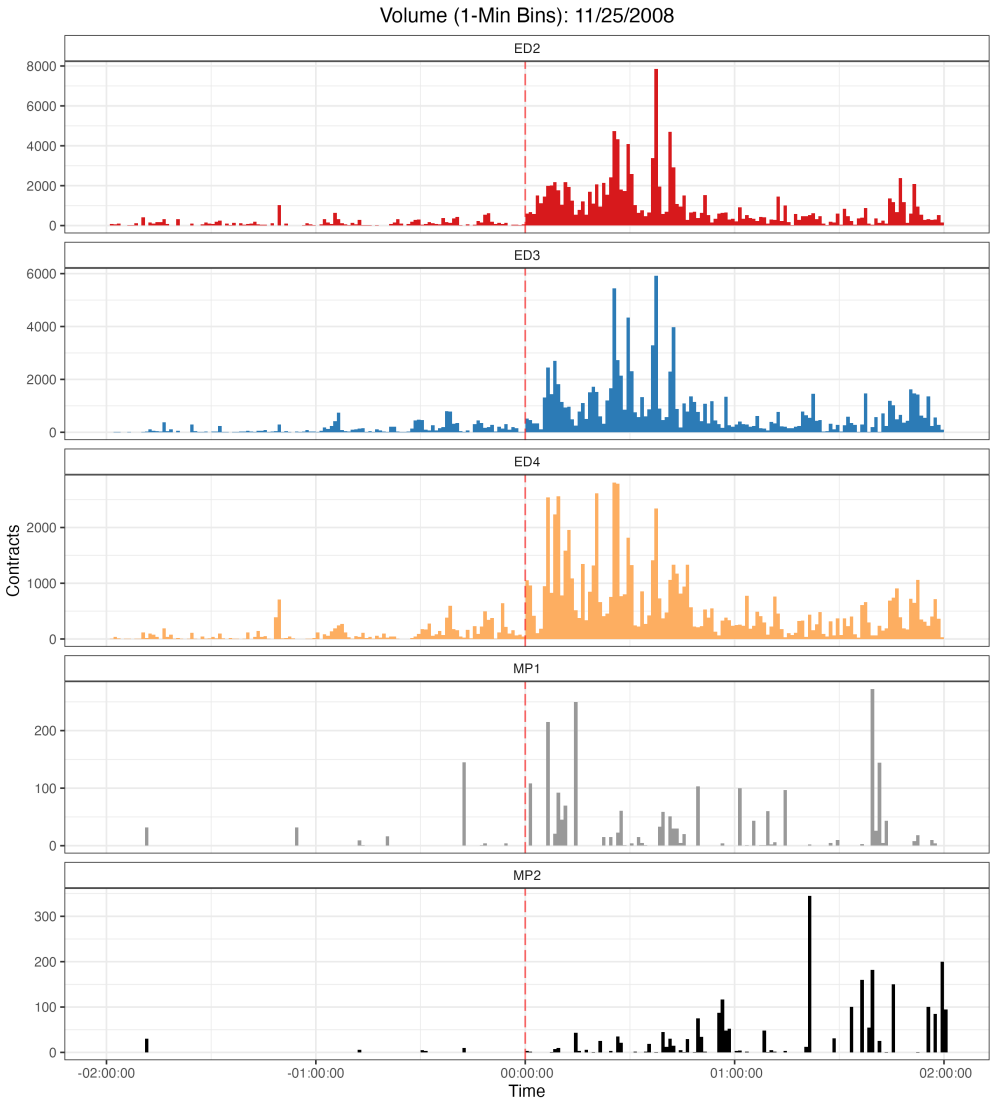


Figure A.6: Trading volume aggregated to 1-minute bins: 11/25/2008

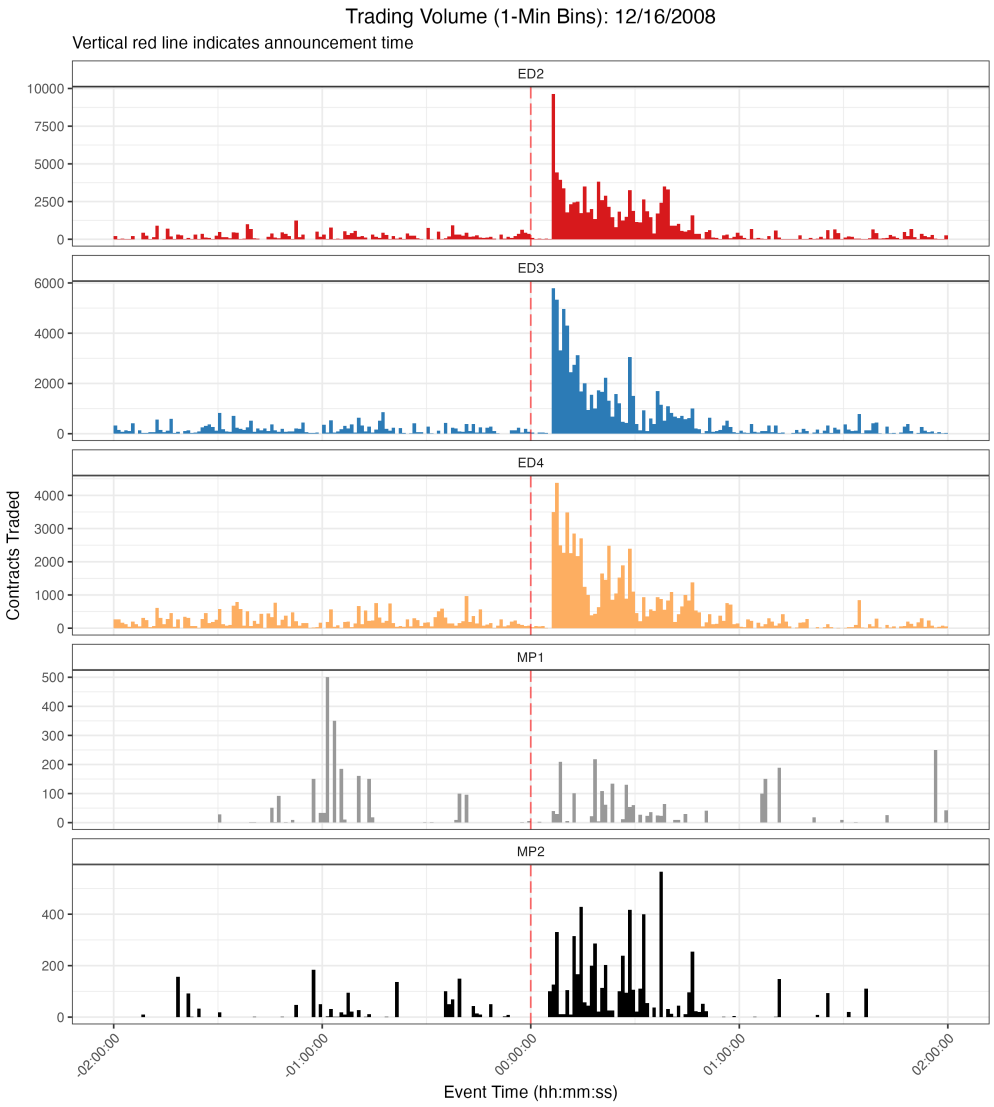


Figure A.7: Trading volume aggregated to 1-minute bins: 12/16/2008

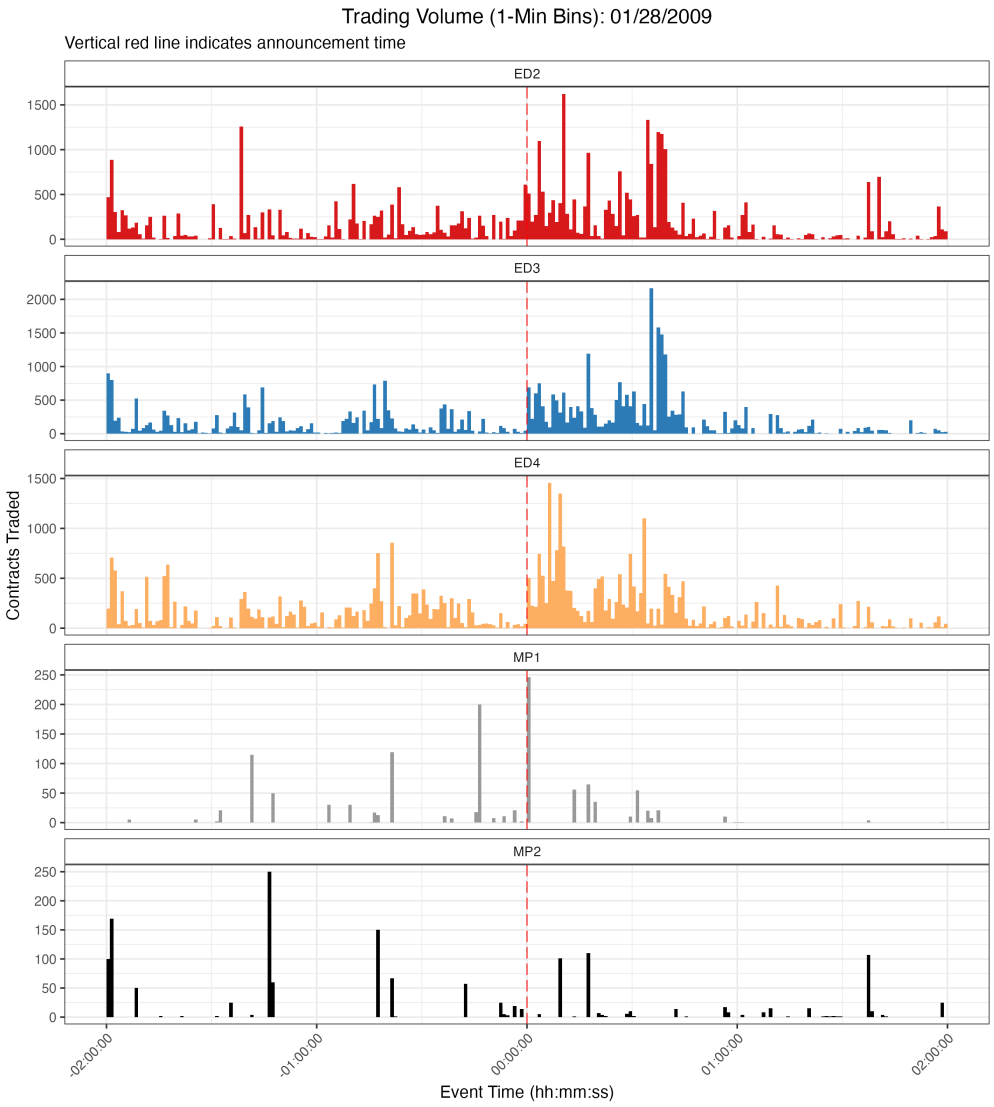


Figure A.8: Trading volume aggregated to 1-minute bins: 01/28/2009

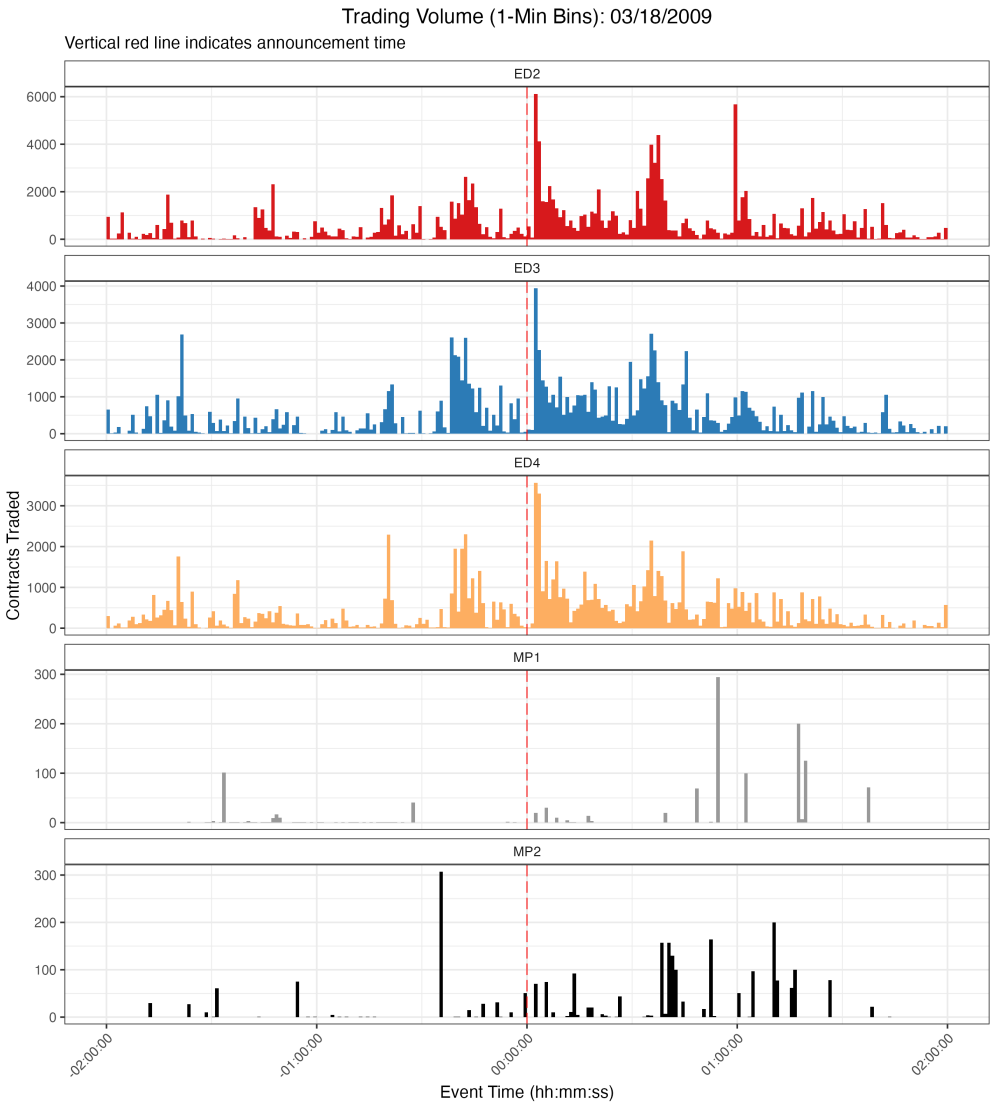


Figure A.9: Trading volume aggregated to 1-minute bins: 03/18/2009

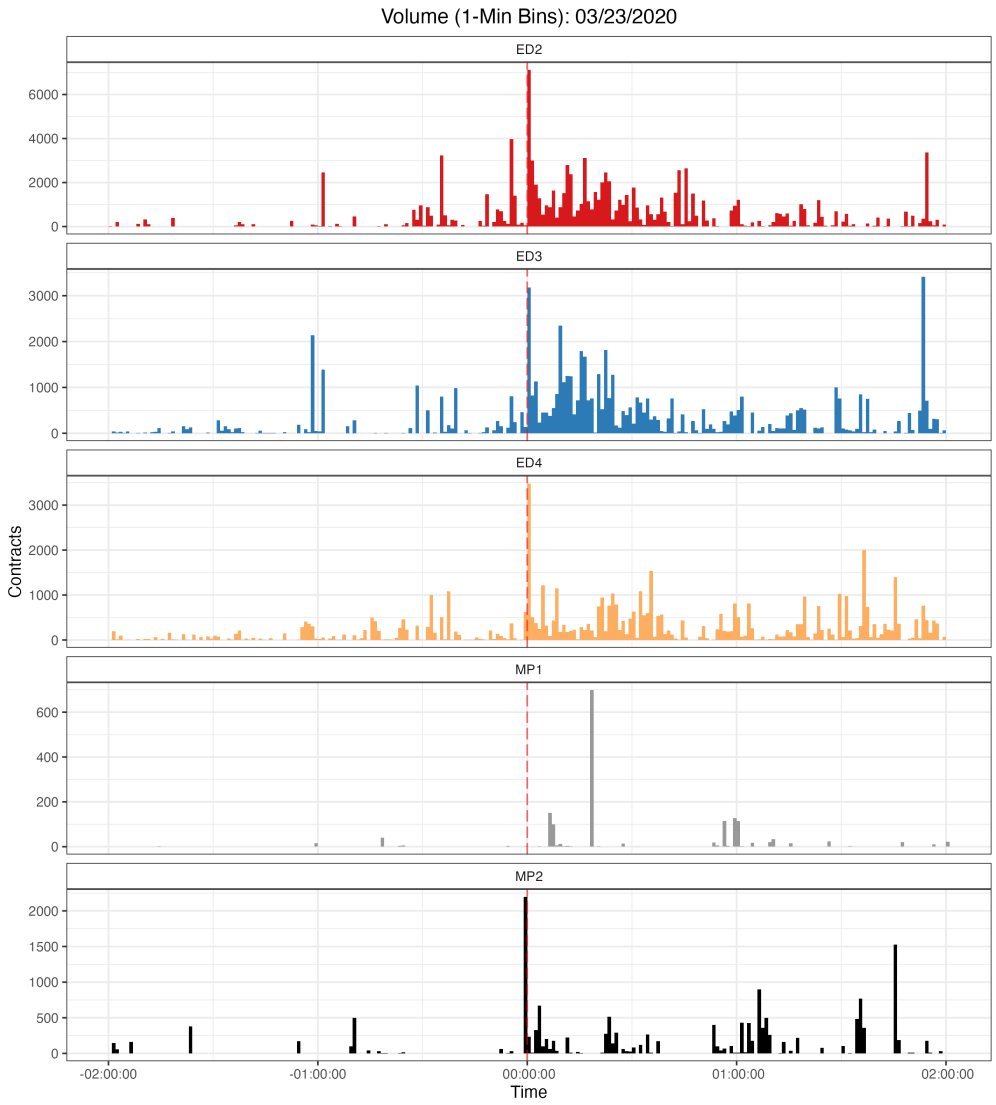


Figure A.10: Trading volume aggregated to 1-minute bins: 03/23/2020

A.3 Market Segmentation

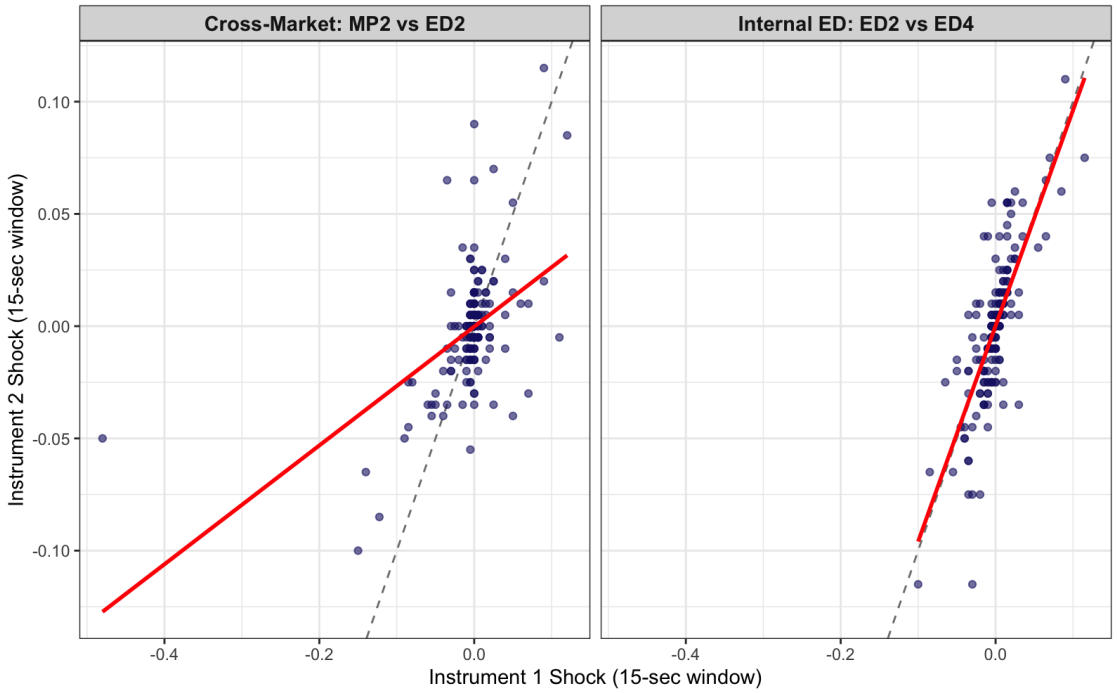


Figure A.11: Market Segmentation: Correlations between MP2 vs. ED2 and ED2 vs. ED4